

US 20040236877A1

(19) **United States**(12) **Patent Application Publication**
Burton(10) **Pub. No.: US 2004/0236877 A1**(43) **Pub. Date: Nov. 25, 2004**

(54) **SWITCH/NETWORK ADAPTER PORT
INCORPORATING SHARED MEMORY
RESOURCES SELECTIVELY ACCESSIBLE
BY A DIRECT EXECUTION LOGIC
ELEMENT AND ONE OR MORE DENSE
LOGIC DEVICES IN A FULLY BUFFERED
DUAL IN-LINE MEMORY MODULE
FORMAT (FB-DIMM)**

application No. 08/992,763, filed on Dec. 17, 1997,
now Pat. No. 6,076,152.

Publication Classification(51) **Int. Cl.⁷** **G06F 13/28**(52) **U.S. Cl.** **710/22**(76) **Inventor: Lee A. Burton, Divide, CO (US)**


Correspondence Address:
HOGAN & HARTSON LLP
ONE TABOR CENTER, SUITE 1500
1200 SEVENTEENTH ST
DENVER, CO 80202 (US)


(21) **Appl. No.: 10/869,199**(22) **Filed: Jun. 16, 2004****Related U.S. Application Data**

(60) Continuation-in-part of application No. 10/618,041, filed on Jul. 11, 2003, which is a continuation-in-part of application No. 10/340,390, filed on Jan. 10, 2003, which is a continuation-in-part of application No. 09/932,330, filed on Aug. 17, 2001, which is a continuation-in-part of application No. 09/755,744, filed on Jan. 5, 2001, now abandoned, which is a division of application No. 09/481,902, filed on Jan. 12, 2000, now Pat. No. 6,247,110, which is a continuation of

(57) **ABSTRACT**

An enhanced switch/network adapter port incorporating shared memory resources ("SNAPM™") selectively accessible by a direct execution logic element and one or more dense logic devices in a fully buffered dual in-line memory module ("FB-DIMM") format for clustered computing systems employing direct execution logic such as multi-adaptive processor elements ("MAP®", all trademarks of SRC Computers, Inc.). Functionally, the SNAPM modules incorporate and properly allocate memory resources so that the memory appears to the associated dense logic device(s) (e.g. a microprocessor) to be functionally like any other system memory such that no time penalties are incurred when accessing it. Through the use of a programmable access coordination mechanism, the control of this memory can be handed off to the SNAPM memory controller and, once in control, the controller can move data between the shared memory resources and the computer network such that the transfer is performed at the maximum rate that the memory devices themselves can sustain. This provides the highest performance link to the other network devices such as MAP® elements, common memory boards and the like.









Fig. 2A






Fig. 3




Fig. 4A




Fig. 4B

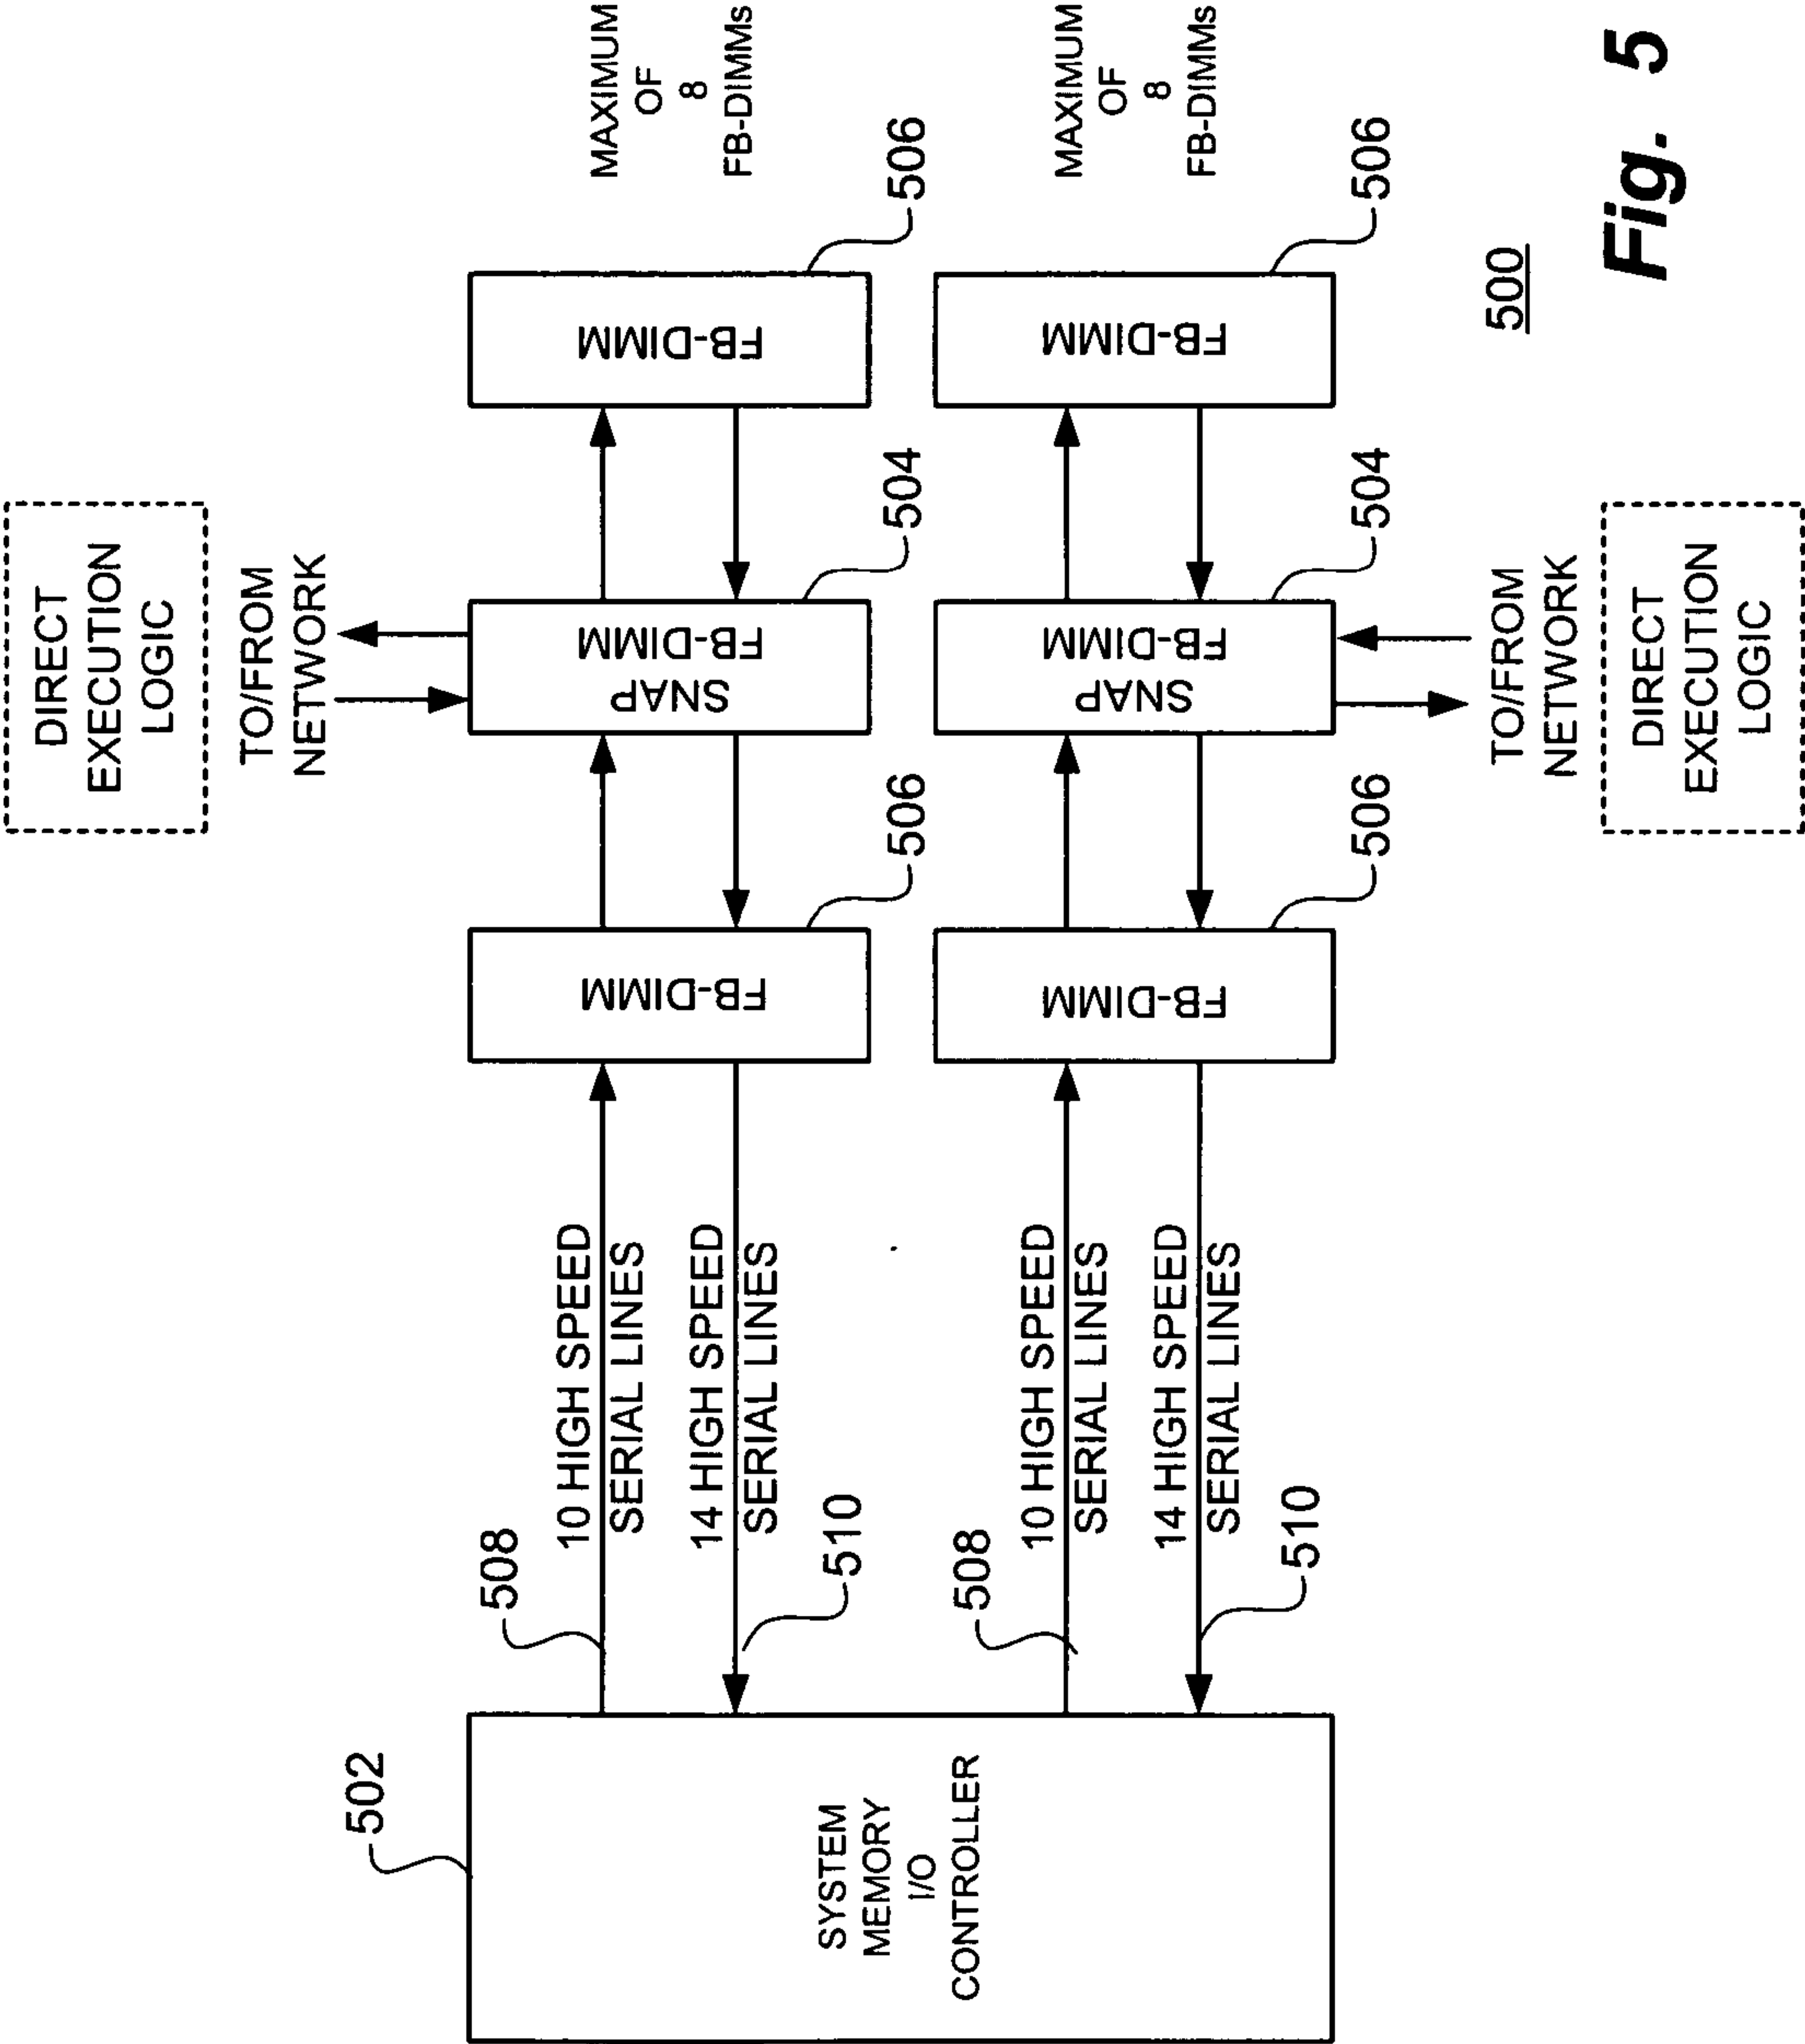


Fig. 5





Fig. 7



504

Fig. 6

**SWITCH/NETWORK ADAPTER PORT
INCORPORATING SHARED MEMORY
RESOURCES SELECTIVELY ACCESSIBLE BY A
DIRECT EXECUTION LOGIC ELEMENT AND
ONE OR MORE DENSE LOGIC DEVICES IN A
FULLY BUFFERED DUAL IN-LINE MEMORY
MODULE FORMAT (FB-DIMM)**

**CROSS REFERENCE TO RELATED PATENT
APPLICATIONS**

[0001] The present invention is a continuation-in-part application and is related to, and claims priority from, U.S. patent application Ser. No. 10/618,041 filed Jul. 11, 2003 for: "Switch/Network Adapter Port Incorporating Shared Memory Resources Selectively Accessible by a Direct Execution Logic Element and One or More Dense Logic Devices", which is a continuation-in-part application and is related to, and claims priority from U.S. patent application Ser. No. 10/340,390 filed Jan. 10, 2003 for: "Switch/Network Adapter Port Coupling a Reconfigurable Processing Element to One or More Microprocessors for Use With Interleaved Memory Controllers, which is a continuation-in-part application and is related to, and claims priority from, U.S. patent application Ser. No. 09/932,330 filed Aug. 17, 2001 for: "Switch/Network Adapter Port for Clustered Computers Employing a Chain of Multi-Adaptive Processors in a Dual In-Line Memory Module Format" which is a continuation-in-part of patent application Ser. No. 09/755,744 filed Jan. 5, 2001 which is a divisional patent application of U.S. patent application Ser. No. 09/481,902 filed Jan. 12, 2000, now U.S. Pat. No. 6,247,110, which is a continuation of Ser. No. 08/992,763, filed Dec. 17, 1997, now U.S. Pat. No. 6,076,152, all of which are assigned to SRC Computers, Inc., Colorado Springs, Colo., the assignee of the present invention, the disclosures of which are herein specifically incorporated in their entirety by this reference.

BACKGROUND OF THE INVENTION

[0002] The present invention relates, in general, to the field of reconfigurable processor-based computing systems. More particularly, the present invention relates to a switch/network adapter port incorporating shared memory resources selectively accessible by a direct execution logic element (such as a reconfigurable computing element comprising one or more field programmable gate arrays "FPGAs") and one or more dense logic devices comprising commercially available microprocessors, digital signal processors ("DSPs"), application specific integrated circuits ("ASICs") and other typically fixed logic components having relatively high clock rates.

[0003] As disclosed in one or more representative embodiments illustrated and described in the aforementioned patents and patent applications, SRC Computers, Inc. proprietary Switch/Network Adapter Port technology (SNAP™, a trademark of SRC Computers, Inc., assignee of the present invention) has previously been enhanced such that the signals from two or more dual in-line memory module ("DIMM") (or Rambus™ in-line memory module "RIMM") slots are routed to a common control chip.

[0004] Physically, in a by-two configuration, two DIMM form factor switch/network adapter port boards may be coupled together using rigid 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 may be utilized in conjunction with more than two DIMM slots, for example, four DIMM slots in a four-way interleaved system.

[0005] In operation, an interleaved memory system may use two or more memory channels running in lock-step wherein a connection is made to one of the DIMM slots and the signals derived are used in conjunction with the original set of switch/network adapter port 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.

[0006] Currently described in the literature is a reconfigurable computing development environment called "Pilchard" which plugs into a personal computer DIMM slot. See, for example, "Pilchard—A Reconfigurable Computing Platform with Memory Slot Interface" developed at the Chinese University of Hong Kong under a then existing license and utilizing SRC Computers, Inc. technology. The Pilchard system, and other present day systems rely on relatively long column address strobe ("CAS") latencies to enable the FPGA to process the memory transactions and are essentially slaves to the memory and I/O controller.

[0007] With the speed gap ever increasing between the processor speeds and the memory subsystem, processor design has been optimized to keep the cache subsystem filled with data that will be needed by the program currently executing on the processor. Thus, the processor itself is becoming less efficient at performing the large block transfers that may be required in certain systems utilizing currently available switch/network devices.

[0008] The need to have a relatively large volume of system dynamic random access memory ("DRAM") has increased in recent years due to the need to handle ever larger databases and with ever increasing problem sizes. At the same time, integrated circuit memory densities continue to double approximately every eighteen to twenty four months. Consequently, more and more memory devices are required in a system to meet an applications needs.

[0009] An even greater impact on the performance of a system has been the ever increasing time (in processor clocks) it takes to access the DRAM in the system. This has created pressure for even faster memory sub-systems. For these reason, the double data rate ("DDR") DDR2 and DDR3 memory specifications have been set forth. These specifications include clock rates of from 200 to 400 MHz, but yet they do not incorporate modifications to the basic interconnect structure and still impose a stub terminated bus structure. Because of the clock rate involved with this bus structure, the number devices present on the bus is limited, thus creating a situation where the memory needs of the applications being run are still not being met.

[0010] For this reason, a new memory bus structure is being developed which is denominated as the Fully Buffered

DIMM (FB-DIMM). The FB-DIMM uses an Advanced Memory Buffer (AMB) to perform serial to parallel conversions necessary to enable the memory controller in the North Bridge to function serially. The Advanced Memory Buffer then converts this to the parallel signaling that is required by the standard DDR2 SRAM. The Advanced Memory Buffer also incorporates a pass-through port to enable the use of multiple FB-DIMM's in a given system. With this bus structure, all of the interconnects are essentially point-to-point differential serial. Further, along with the pass-through port, a vary large memory subsystem can be created.

SUMMARY OF THE INVENTION

[0011] In order to increase processor operational efficiency in conjunction with a switch/network adapter port, the present invention advantageously incorporates and properly allocates memory resources, such as dynamic random access memory ("DRAM"), located on the module itself. Functionally, this memory appears to the dense logic device (e.g. a microprocessor) to be like other system memory and no time penalties are incurred when reading to, or writing from, it.

[0012] Through the use of an access coordination mechanism, the control of this memory can be handed off to the switch/network adapter port memory controller. Once in control, the controller can move data between the memory resources and the computer network, based for example, on control parameters that may be located in on-board registers. This data movement is performed at the maximum rate that the memory devices themselves can sustain, thereby providing the highest performance link to the other network devices such as direct execution logic devices such as Multi-Adaptive Processing elements (MAP® a trademark of SRC Computers, Inc.), common memory boards and the like.

[0013] Unlike the Pilchard system described previously, the system and method of the present invention does not need to rely on relatively long CAS memory latencies to enable the associated FPGA to process the memory transactions. Moreover, the system and method of the present invention functions as a true peer to the system memory and I/O controller and access to the shared memory resources is arbitrated for between the memory and I/O controller and the switch/network adapter port controller.

[0014] Further, with increasing system security demands, as well as other functions that require unique memory address access patterns, the addition of a programmable memory controller to the system/network adapter port control unit enables this improved system to meet these needs. Functionally, the memory controller is enabled such that the address access patterns utilized in the performance of the data movement to and from the collocated memory resources is programmable. This serves to effectively eliminate the performance penalty that is common when performing scatter/gather and other similar functions.

[0015] In a representative embodiment of the present invention disclosed herein, the memory and I/O controller, as well as the enhanced switch/network adapter port memory ("SNAPM™") controller, can control the common memory resources on the SNAPM modules through the inclusion of various data and address switches (e.g. field effect transistors "FETs", or the like) and tri-stable latches.

These switching resources and latches are configured such that the data and address lines may be driven by either the memory and I/O controller or the SNAPM memory controller while complete DIMM (and RIMM or other memory module format) functionality is maintained. Specifically, this may be implemented in various ways including the inclusion of a number of control registers added to the address space accessible by the memory and I/O controller which are used to coordinate the use of the shared memory resources.

[0016] In operation, when the memory and I/O controller is in control, the SNAPM memory controller is barred from accessing the DRAM memory. Conversely, when the SNAPM memory controller is in control, the address/control and data buses from the memory and I/O controller are disconnected from the DRAM memory. However, the SNAPM memory controller continues to monitor the address and control bus for time critical commands such as memory refresh commands. Should the memory and I/O controller issue a refresh command while the SNAPM memory controller is in control of the DRAM memory, it will interleave the refresh command into its normal command sequence to the DRAM devices. Additionally, when the memory and I/O controller is in control, the SNAPM modules monitor the address and command bus for accesses to any control registers located on the module and can accept or drive replies to these commands without switching control of the collocated memory resources.

[0017] Functionally, the SNAPM controller contains a programmable direct memory access ("DMA") engine which can perform random access and other DMA operations based on the state of any control registers or in accordance with other programmable information. The SNAPM controller is also capable of performing data re-ordering functions wherein the contents of the DRAM memory can be read out and then rewritten in a different sequence.

[0018] Particularly disclosed herein is a computer system comprising at least one dense logic device, a controller for coupling the dense logic device to a control block and a memory bus, one or more memory module slots coupled to the memory bus with at least one of the memory module slots comprising a buffered memory module, an adapter port including shared memory resources associated with a subset of the plurality of memory module slots and a direct execution logic element coupled to the adapter port. The dense logic device and the direct execution logic element may both access the shared memory resources. In a preferred embodiment, the adapter port may be conveniently provided in an FB-DIMM, or other buffered memory module form factor.

[0019] Also disclosed herein is a computer system comprising at least one dense logic device, an interleaved controller for coupling the dense logic device to a control block and a memory bus, a plurality of memory slots coupled to the memory bus with at least one of the memory module slots comprising a buffered memory module, an adapter port including shared memory resources associated with at least two of the memory slots and a direct execution logic element coupled to at least one of the adapter ports.

[0020] Further disclosed herein is a computer system including an adapter port for electrical coupling between a memory bus of the computer system and a network inter-

face. The computer system comprises at least one dense logic device coupled to the memory bus and the adapter port comprises a memory resource associated with the adapter port and a control block for selectively enabling access by the dense logic device to the memory resource. In a particular embodiment disclosed herein, the computer system may further comprise an additional adapter port having an additional memory resource associated with it and the control block being further operative to selectively enable access by the dense logic device to the additional memory resource.

[0021] Broadly, the system and method of the present invention disclosed herein includes a switch/network adapter port with collocated memory in an FB-DIMM format that may be isolated to allow peer access to the memory by either a system memory and I/O controller or switch/network adapter port memory controller. The switch/network adapter port with on-board memory disclosed may be utilized as an interface itself and also allows the switch/network adapter port memory controller to operate directly on data retained in the shared memory resources. This enables it to prepare the data for transmission in operations requiring access to a large block of non-sequential data, such as scatter and gather. The system and method of the present invention described herein further discloses a switch/network adapter port with shared memory resources which incorporates a smart, fully parameterized DMA engine providing the capability of performing scatter/gather and other similar functions.

BRIEF DESCRIPTION OF THE DRAWINGS

[0022] 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 description of a preferred embodiment taken in conjunction with the accompanying drawings, wherein:

[0023] **FIG. 1** is a functional block diagram of a switch/network adapter port for a clustered computing system employing a chain of multi-adaptive processors in a DIMM format functioning as direct execution logic to significantly enhance data transfer rates over that otherwise available from the peripheral component interconnect ("PCI") bus;

[0024] **FIG. 2A** is a functional block diagram of an exemplary embodiment of a switch/network adapter port incorporating collocated shared memory resources illustrating in a by-two configuration of interleaved DIMM slot form factor SNAPM elements coupled to a common SNAPM memory control element for coupling to a cluster interconnect fabric including one or more direct execution logic devices such as MAP® elements;

[0025] **FIG. 2B** is a further functional block diagram of another exemplary embodiment of a switch/network adapter port incorporating collocated shared memory resources in accordance with the present invention illustrating a by-four configuration of interleaved DIMM slot form factor SNAPM elements coupled to a common SNAPM memory control element;

[0026] **FIG. 3** is a functional block diagram of a representative embodiment of a by-two SNAPM system in accordance with the present invention comprising a pair of circuit

boards, each of which may be physically and electrically coupled into one of two DIMM memory slots, and one of which may contain a SNAPM control block in the form of a field programmable gate array ("FPGA") functioning as the SNAPM memory control block of the preceding **FIGS. 2A and 2B**;

[0027] **FIG. 4A** is a corresponding functional block diagram of the embodiment of the preceding figure wherein the memory and I/O controller drives the address/control and data buses for access to the shared memory resources of the SNAPM elements through the respective address and data switches;

[0028] **FIG. 4B** is an accompanying functional block diagram of the embodiment of **FIG. 3** wherein the SNAPM memory control block provides access to the shared memory resources and disconnects the address/control and data buses from the system memory and I/O controller;

[0029] **FIG. 5** is a functional block diagram of a representative fully buffered DIMM memory system implemented in accordance with a particular embodiment of the present invention and wherein the number of sets of FB-DIMM branches is based on the bandwidth requirements of the system;

[0030] **FIG. 6** is a corresponding functional block diagram of a representative switch/network adapter port FB-DIMM block for possible use in conjunction with the FB-DIMM memory system of the preceding figure and wherein the double data rate synchronous dynamic random access memory (DDR SDRAM) array in this exemplary embodiment is shown as being 72 bits wide; and

[0031] **FIG. 7** is a simplified view of a typical FB-DIMM memory module which is also coupled to the memory controller system maintenance (SM) bus and a clock (CLK) signal source.

DESCRIPTION OF A REPRESENTATIVE EMBODIMENT

[0032] With reference now to **FIG. 1**, a functional block diagram of an 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 functioning as direct execution logic elements in a DIMM format to significantly enhance data transfer rates over that otherwise available from the peripheral component interconnect ("PCI") bus.

[0033] In the particular embodiment illustrated, the computer system **100** includes one or more dense logic devices in the form of processors **102₀** and **102₁** which are coupled to an associated memory and I/O controller **104** (e.g. a "North Bridge"). In the operation of the particular embodiment illustrated, the controller **104** sends and receives control information from a separate PCI control block **106**. It should be noted, however, that in alternative implementations of the present invention, the controller **104** and/or the PCI control block **106** (or equivalent) may be integrated within the processors **102** themselves and that 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 alternative embodiments, the card slots **108** may alternatively comprise PCI-X, PCI Express, accelerated graphics port (“AGP”) or system maintenance (“SM”) bus connections.

[0034] 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 system shown, a DIMM MAP® element **112** may be associated with, or physically located within, one of the DIMM slots **114**. Control information to or from the DIMM MAP® element **112** may be 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 interconnect fabric connection **120** connected to MAP® chain ports. It should be noted that, the DIMM MAP® element **112** may also comprise a Rambus™ DIMM (“RIMM”) MAP® element.

[0035] Since the DIMM memory located within the DIMM slots **114** comprises the primary storage location for the microprocessor(s) **102₀**, **102₁**, it is 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 be realized as disclosed in the aforementioned patents and patent applications.

[0036] To this end, by placing the DIMM MAP® element **112** in one of the PC’s DIMM slots **114**, its control chip can accept the normal memory “read” and “write” transactions and convert them to a format used by an interconnect switch or network. To this end, each MAP® element **112** may also include 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 application, the DIMM MAP® element **112** would extract the data from the packet and store it until needed by the receiving processor **102**.

[0037] 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 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 all other bus agents.

[0038] 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 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 possible implementation, the system maintenance (“SM”) bus (not shown) could also be 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**.

[0039] With a DIMM MAP® element **112** associated with what might be an entire DIMM slot **114**, the system 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 Intel™ chip sets used in most personal computers (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 consumed by this transfer.

[0040] It should also be noted that in certain computer systems, 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 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 the interconnect fabric.

[0041] With reference additionally now to FIG. 2A, a functional block diagram of an exemplary embodiment of a switch/network adapter port **200A** incorporating collocated common memory resources in accordance with the present invention is shown. In this regard, like structure and functionality to that disclosed with respect to the foregoing figure is here like numbered and the foregoing description thereof shall suffice herefor. The switch/network adapter port with common memory (“SNAPM”) **200A** is shown in an exemplary by-two configuration of interleaved DIMM slot form factor SNAPM elements **204** (SNAPM A and SNAPM B) each coupled to a common control element **202** (comprising, together with the two SNAPM elements **204** “SNAPM”) and with each of the SNAPM elements **204** including respective DRAM memory **206A** and **206B** in conjunction with associated switches and buses **208A** and **208B** respectively as will be more fully described hereinafter. In this embodiment, the controller **104** is an interleaved memory controller bi-directionally coupled to the DIMM slots **114** and SNAPM elements **204** by means of a Channel A **216A** and a Channel B **216B**.

[0042] With reference additionally now to **FIG. 2B**, a functional block diagram of another exemplary embodiment of a switch/network adapter port **200B** incorporating collocated common memory resources in accordance with the present invention is shown. Again, like structure and 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** with common memory is shown in a by-four configuration of interleaved DIMM slot form factor SNAPM elements **204** coupled to a common SNAPM memory control element **202** (comprising, together with the four SNAPM elements **204** "SNAPM"). In this embodiment, the controller **104** is again an interleaved memory controller bi-directionally coupled to the DIMM slots **114** and SNAPM elements **204** by means of a respective Channel A **216A**, Channel B **216B**, Channel C **216C** and Channel D **216D**.

[0043] With reference additionally now to **FIG. 3**, a functional block diagram of a representative embodiment of a by-two SNAPM system **300** in accordance with the present invention is shown. The SNAPM system, in the exemplary embodiment shown, comprises a pair of circuit boards **204**, each of which may be physically and electrically coupled into one of two DIMM (RIMM or other memory module form factor) memory slots, and one of which may contain a SNAPM control block **202** in the form of, for example, an FPGA programmed to function as the SNAPM memory control block of the preceding **FIGS. 2A and 2B**.

[0044] Each of the SNAPM circuit boards **204** comprises respective collocated common memory resources **206A** ("Memory A") and **206B** ("Memory B") which may be conveniently provided in the form of DRAM, SRAM or other suitable memory technology types. Each of the memory resources **206A** and **206B** is respectively associated with additional circuitry **208A** and **208B** comprising, in pertinent part, respective DIMM connectors **302A** and **302B**, a number of address switches **304A** and **304B** and a number of data switches **306A** and **306B** along with associated address/control and data buses. The address switches **304A** and **304B** and data switches **306A** and **306B** are controlled by a switch direction control signal provided by the SNAPM control block **202** on control line **308** as shown. The address switches **304** and data switches **306** may be conveniently provided as FETs, bipolar transistors or other suitable switching devices. The network connections **120** may be furnished, for example, as a flex connector and corresponds to the cluster interconnect fabric of the preceding figures for coupling to one or more elements of direct execution logic such as MAP® elements available from SRC Computers, Inc.

[0045] With reference additionally now to **FIG. 4A**, a corresponding functional block diagram of the embodiment of the preceding figure is shown wherein the memory and I/O controller (element **104** of **FIGS. 1, 2A and 2B**) drives the address/control and data buses for access to the shared memory resources **206** of the SNAPM elements **204** through the respective address and data switches **304** AND **306** in accordance with the state of the switch direction control signal on control line **308**.

[0046] With reference additionally now to **FIG. 4B**, an accompanying functional block diagram of the embodiment of **FIG. 3** is shown wherein the SNAPM memory control

block **202** provides access to the shared memory resources **206** and disconnects the address/control and data buses from the system memory and I/O controller (element **104** of **FIGS. 1, 2A and 2B**) in accordance with an opposite state of the switch direction control signal on control line **308**.

[0047] As shown with respect to **FIGS. 4A and 4B**, the memory and I/O controller (element **104** of **FIGS. 1, 2A and 2B**), as well as the SNAPM memory controller **202**, can control the common memory resources **206** on the SNAPM modules **204**. The switches **304** and **306** are configured such that the data and address lines may be driven by either the memory and I/O controller **104** or the SNAPM memory controller **202** while complete DIMM (and RIMM or other memory module format) functionality is maintained. Specifically, this may be implemented in various ways including the inclusion of a number of control registers added to the address space accessible by the memory and I/O controller **104** which are used to coordinate the use of the shared memory resources **206**. In the embodiment illustrated, the least significant bit ("LSB") data lines (07:00) of lines (71:00) and/or selected address bits may be used to control the SNAPM control block **202**, and hence, the allocation and use of the shared memory resources **206**.

[0048] In operation, when the memory and I/O controller **104** is in control, the SNAPM memory controller **202** is barred from accessing the DRAM memory **206**. Conversely, when the SNAPM memory controller **202** is in control, the address/control and data buses from the memory and I/O controller **104** are disconnected from the DRAM memory **206**. However, the SNAPM memory controller **202** continues to monitor the address and control bus for time critical commands such as memory refresh commands. Should the memory and I/O controller **104** issue a refresh command while the SNAPM memory controller **202** is in control of the DRAM memory **206**, it will interleave the refresh command into its normal command sequence to the DRAM devices. Additionally, when the memory and I/O controller **104** is in control, the SNAPM modules **204** monitor the address and command bus for accesses to any control registers located on the module and can accept or drive replies to these commands without switching control of the collocated memory resources **206**.

[0049] With reference additionally now to **FIG. 5**, a functional block diagram of a representative fully buffered DIMM memory system **500** implemented in accordance with a particular embodiment of the present invention is shown wherein the number of sets of FB-DIMM branches is based on the bandwidth requirements of the system.

[0050] The FB-DIMM memory system **500** comprises, in pertinent part a system memory I/O controller **502** which is analogous to the memory and I/O controller **104** of the preceding figures. One or more switch/network adapter port FB-DIMM blocks **504**, which will be described in more detail hereinafter, may be physically and electrically coupled to standard DIMM slots within the memory system **500** and are bidirectionally coupled to a computer network comprising one or more direct execution logic blocks as shown. In like manner, a number of FB-DIMM memory modules **506** are also physically and electrically coupled to standard DIMM slots within the memory system **500** and, in the representative embodiment shown, a maximum of eight FB-DIMM modules may be provided. The switch/network

adapter port FB-DIMM blocks **504** and the FB-DIMM memory modules **506** are coupled to the system memory I/O controller **502** through ten high speed serial lines **508** and fourteen high speed serial lines **510** as illustrated.

[0051] With reference additionally now to **FIG. 6**, a corresponding functional block diagram of a representative switch/network adapter port FB-DIMM block **504** is shown for possible use in conjunction with the FB-DIMM memory system **500** of the preceding figure. The switch/network adapter port FB-DIMM block **504** comprises, in pertinent part, a SNAP Advanced Memory Buffer (“AMB”) control FPGA analogous to the SNAPM control block **202** of the preceding figures which provides the bi-directional coupling to the direct execution logic of the computer network. It further includes a number of double data rate two synchronous dynamic random access memory (DDR2 SDRAM) elements **604** in an array which, in this exemplary embodiment, is shown as being 72 bits wide. The SNAP AMB control FPGA **602** is coupled to the DDR2 SDRAM elements **604** through an address/control (“ADR/CTL”) bus **606** and a bidirectional data bus **608**. The SNAP AMB control FPGA **602** of the switch/network adapter port FB-DIMM block **504** is electrically (and physically) coupled to a FB-DIMM connector **610** through a pair of high speed serial lines **508** and fourteen high speed serial lines **510**, one pair of which function as high speed pass-through serial lines as illustrated.

[0052] As previously illustrated in the embodiments of **FIG. 3**, the address switches **304** and data switches **306** of a SNAPM system **300** may be conveniently provided as FETs, bipolar transistors or other suitable switching devices to provide isolation between the SNAP control FPGA and the North Bridge. With an FB-DIMM based system as herein disclosed, the Advanced Memory Buffer naturally provides an analogous isolation point. Therefore, by constructing the AMB out of the SNAP FPGA devices, the functionality of the SNAPM system **300** can be effectively duplicated with a significant reduction in the complexity of the overall module design. In this particular embodiment, the address switches **304** and data switches **306** of the SNAPM system **300** are no longer required because of the conversion necessary to go from the serial format to the parallel format of the SDRAMs. The pass through port allows the SNAP FB-DIMM block **504** to ‘claim’ the transaction or pass it on to the next FB-DIMM memory module **506** based on the address of the transaction.

[0053] The SNAP control FPGA is capable of providing all of the specified AMB functionality. Additionally, the SNAP controller may be conveniently configured to provide control registers that can enable the local SDRAM to be exclusively controlled by SNAP. By utilizing the pass through port, normal system memory traffic can still occur and future clock increases, either in the serial interface or in the SDRAM components, would more easily be accommodated. Further, to the extent the AMB is, or in the future is, configurable through downloadable parameters, the FB-DIMM block of the present invention could likewise be reprogrammable in the way the associated memory may be accessed.

[0054] With reference additionally now to **FIG. 7**, a simplified view of a typical FB-DIMM memory module **506** is shown in a card **702** form factor for electrical and physical

retention within a standard DIMM memory slot. The FB-DIMM memory module **506** comprises an on card buffer **704** as well as a number of DRAM elements **706** all electrically accessible through a card edge connector **708**. As illustrated, the ten high speed serial lines **508** and fourteen high speed serial lines **510** are coupled to the buffer **704** through the edge connector **708** as is the system memory I/O controller **502** system maintenance (SM) bus **712** and a clock (CLK) signal source **710**.

[0055] While there have been described above the principles of the present invention in conjunction with specific module configurations and circuitry, 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 application or of any further application derived therefrom.

What is claimed is:

1. A computer system comprising:
 - at least one dense logic device;
 - a controller for coupling said at least one dense logic device to a control block and a memory bus;
 - one or more memory module slots coupled to said memory bus, at least one of said one or more memory module slots comprising a buffered memory module;
 - an adapter port associated with a subset of said one or more memory module slots, said adapter port including associated memory resources; and
 - at least one direct execution logic element coupled to said adapter port, said memory resources being selectively accessible by said at least one dense logic device and said at least one direct execution logic element.
2. The computer system of claim 1 wherein said controller comprises an interleaved memory controller.
3. The computer system of claim 1 wherein said buffered memory module comprises an FB-DIMM memory module.
4. The computer system of claim 3, wherein said adapter port comprises an FB-DIMM physical format for retention within one of said memory module slots.
5. The computer system of claim 1 wherein said control block provides control information to said adapter port.
6. The computer system of claim 1 wherein said control block provides control information to said direct execution logic element.

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

8. The computer system of claim 7 wherein said peripheral bus control block provides control information to said adapter port.

9. The computer system of claim 7 wherein said peripheral control block provides control information to said direct execution logic element.

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

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

12. The computer system of claim 10 wherein said graphics control block provides control information to said direct execution logic element.

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

14. The computer system of claim 13 wherein said systems maintenance control block provides control information to said adapter port.

15. The computer system of claim 13 wherein said systems maintenance control block provides control information to said direct execution logic element.

16. The computer system of claim 1 wherein said direct execution logic element comprises a reconfigurable processor element.

17. The computer system of claim 1 wherein said direct execution logic element is operative to alter data received from said controller on said memory bus.

18. The computer system of claim 1 wherein said direct execution logic element is operative to alter data received from an external source prior to placing altered data on said memory bus.

19. The computer system of claim 1 wherein said direct execution logic element comprises:

a control block coupled to said adapter port.

20. The computer system of claim 19 wherein said direct execution logic 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.

21. The computer system of claim 20 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.

22. The computer system of claim 1 wherein said direct execution logic element comprises:

a chain port for coupling said direct execution logic element to another direct execution logic element.

23. The computer system of claim 19 wherein said direct execution logic element further comprises:

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

24. A computer system comprising:

at least one dense logic device;

an interleaved controller for coupling said at least one dense logic device to a control block and a memory bus;

a plurality of memory slots coupled to said memory bus, at least one of said plurality of memory slots comprising a buffered memory module;

an adapter port associated with at least two of said plurality of memory slots, each of said adapter port including associated memory resources; and

a direct execution logic element coupled to at least one of said adapter ports, said memory resources being selectively accessible by said at least one dense logic device and said direct execution logic element.

25. The computer system of claim 24 wherein said plurality of memory slots comprise FB-DIMM memory module slots.

26. The computer system of claim 25 wherein said adapter port comprises an FB-DIMM physical format for retention within one of said FB-DIMM memory module slots.

27. The computer system of claim 24 wherein said control block provides control information to said adapter port.

28. The computer system of claim 24 wherein said control block provides control information to said direct execution logic element.

29. The computer system of claim 24 wherein said control block comprises a peripheral bus control block.

30. The computer system of claim 29 wherein said peripheral bus control block provides control information to said adapter port.

31. The computer system of claim 29 wherein said peripheral control block provides control information to said direct execution logic element.

32. The computer system of claim 24 wherein said control block comprises a graphics control block.

33. The computer system of claim 32 wherein said graphics control block provides control information to said adapter port.

34. The computer system of claim 32 wherein said graphics control block provides control information to said direct execution logic element.

35. The computer system of claim 24 wherein said control block comprises a systems maintenance control block.

36. The computer system of claim 35 wherein said systems maintenance control block provides control information to said adapter port.

37. The computer system of claim 35 wherein said systems maintenance control block provides control information to said direct execution logic element.

38. The computer system of claim 24 wherein said control block comprises a PCI-X control block.

39. The computer system of claim 38 wherein said PCI-X control block provides control information to said adapter port.

40. The computer system of claim 38 wherein said PCI-X control block provides control information to said direct execution logic element.

41. The computer system of claim 24 wherein said control block comprises a PCI Express control block.

42. The computer system of claim 41 wherein said PCI Express control block provides control information to said adapter port.

43. The computer system of claim 41 wherein said PCI Express control block provides control information to said direct execution logic element.

44. The computer system of claim 24 wherein said direct execution logic element comprises a reconfigurable processor element.

45. The computer system of claim 24 wherein said direct execution logic element is operative to alter data received from said controller on said memory bus.

46. The computer system of claim 24 wherein said direct execution logic element is operative to alter data received from an external source prior to placing altered data on said memory bus.

47. The computer system of claim 24 wherein said direct execution logic element comprises:

a control block coupled to said adapter port.

48. The computer system of claim 47 wherein said direct execution logic 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.

49. The computer system of claim 48 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.

50. The computer system of claim 24 wherein said direct execution logic element comprises:

a chain port for coupling said processor element to another direct execution logic element.

51. The computer system of claim 47 wherein said direct execution logic element further comprises:

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

52. A computer system including an adapter port for electrical coupling between a memory bus of said computer system and a network interface, said computer system comprising at least one dense logic device coupled to said memory bus through a memory module connector, said adapter port comprising:

a memory resource associated with said adapter port; and

a control block for selectively enabling access by said at least one dense logic device to said memory resource.

53. The computer system of claim 52 wherein said control block is further operational to selectively preclude access by said at least one dense logic device to said memory resource.

54. The computer system of claim 52 further comprising:

at least one direct execution logic element coupled to said network interface.

55. The computer system of claim 54 wherein said control block is further operational to alternatively enable access to said memory resource by said at least one dense logic device and said at least one direct execution logic element.

56. The computer system of claim 52 wherein said memory bus further comprises at least one memory module slot and said adapter port is configured for physical retention within said at least one memory module slot through said memory module connector.

57. The computer system of claim 56 wherein said at least one memory module slot comprises an FB-DIMM slot.

58. The computer system of claim 52 further comprising:

an additional adapter port;

an additional memory resource associated with said additional adapter port, said control block further operative

to selectively enable access by said at least one dense logic device to said additional memory resource.

59. The computer system of claim 58 wherein said control block is further operational to selectively preclude access by said at least one dense logic device to said memory resource and said additional memory resource.

60. The computer system of claim 59 further comprising at least one direct execution logic element coupled to said network interface.

61. The computer system of claim 60 wherein said control block is further operational to alternatively enable access to said memory resource and said additional memory resource by said at least one dense logic device and said at least one direct execution logic element.

62. The computer system of claim 58 wherein said memory bus further comprises first and second memory module slots for physical retention of said at least one adapter port and said additional adapter port respectively.

63. The computer system of claim 62 wherein said first and second memory module slots comprise FB-DIMM slots.

64. The computer system of claim 58 wherein said control block is located on a module comprising said adapter port.

65. The computer system of claim 52 wherein said computer system further comprises:

a memory and I/O controller interposed between said at least one dense logic device and said memory bus.

66. The computer system of claim 65 wherein said memory and I/O controller comprises an interleaved memory controller.

67. The computer system of claim 52 wherein said memory bus comprises address/control and data portions thereof.

68. The computer system of claim 52 wherein said memory bus provides address/control and data inputs to said control block to at least partially control its functionality.

69. The computer system of claim 52 wherein said control block further comprises a DMA controller for providing direct memory access operations to said memory resource.

70. The computer system of claim 69 wherein said DMA controller is fully parameterized.

71. The computer system of claim 69 wherein said DMA controller enables scatter/gather functions to be implemented.

72. The computer system of claim 69 wherein said DMA controller enables irregular data access pattern functions to be implemented.

73. The computer system of claim 69 wherein said DMA controller enables data packing functions to be implemented.

74. The computer system of claim 52 wherein said memory resource may be isolated from said memory bus in response to said control block to enable access thereto by a device coupled to said network interface.

75. The computer system of claim 52 wherein said memory resource comprises random access memory.

76. The computer system of claim 75 wherein said random access memory comprises DRAM.