




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 13/40	A1	(11) International Publication Number: WO 97/44739 (43) International Publication Date: 27 November 1997 (27.11.97)
(21) International Application Number: PCT/US97/00891 (22) International Filing Date: 23 January 1997 (23.01.97) (30) Priority Data: 08/652,787 23 May 1996 (23.05.96) US (71) Applicant: ADVANCED MICRO DEVICES, INC. [US/US]; 5204 East Ben White Boulevard, Mail Stop 562, Austin, TX 78741 (US). (72) Inventors: LEE, Sherman; 28531 Cedarbluff Drive, Rancho Palos Verdes, CA 90275 (US). HALLIGAN, JoAnne, K.; 9109 Edwardson Lane, Austin, TX 78749 (US). (74) Agent: DRAKE, Paul, S.; Advanced Micro Devices, Inc., 5204 East Ben White Boulevard, M/S 562, Austin, TX 78741 (US).		(81) Designated States: JP, KR, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i>

(54) Title: APPARATUS FOR CONVERTING DATA BETWEEN DIFFERENT ENDIAN FORMATS AND SYSTEM AND METHOD EMPLOYING SAME

(57) Abstract

A byte swapping device includes first and second data ports and data path logic coupled between the first and second data ports. The byte swapping device is employed in a data processing system comprising a data storage device configured to store bytes of data, a processor which reads data from the data storage device and writes data to the data storage device, and the bytes wapping device coupled between the data storage device and the processor. The first data port is coupled to the storage device and the second data port is coupled to the processor. The storage device is typically a system memory or peripheral device controller. The processor processes data in a first endian format, i.e., big-endian or little-endian format, and at least a portion of the data stored in the data storage device is in the opposite byte ordering. The byte swapping device selectively byte swaps data transferred between the processor and storage device. In the preferred embodiment, data conversion apertures, or ranges, are defined in the processor address space and the processor provides address signals to the byte swapping device. The byte swapping device selectively byte swaps the data based upon the relationship between the addresses received by the byte swapping device and the data conversion apertures. In one embodiment, the processor programs aperture storage elements with the values of the data conversion apertures. In another embodiment, the data conversion apertures are fixed. In an alternate embodiment, the processor provides control signals to the byte swapping device, wherein the byte swapping device selectively converts the data in response to the control signals from the processor. In one embodiment, the processor is configured to execute a characteristic instruction set, wherein the processor provides the one or more control signals to the byte swapping device in response to which instruction in the instruction set the processor executes.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

Title: Apparatus for Converting Data Between Different Endian Formats and System and Method Employing Same

Field of the Invention

5 The present invention relates generally to digital systems with data of different endian formats, and more particularly to the conversion between little-endian and big-endian formats.

Description of the Related Art

10 A typical computer system includes among its components a central processing unit (CPU) coupled to a system memory by an address bus, a data bus and a control bus. The CPU has an associated address range which comprises the addresses which the CPU may supply on the address bus. This address range is also referred to as the address space of the CPU or system. The system memory typically occupies a portion of the address space of the system. Other devices, such as storage devices, display devices, or other input/output devices also occupy a portion of the address space. The system memory is arranged as a large array of cells,
15 each cell having an associated memory address in the system address space. In most computers a cell is a byte, a byte being eight binary digits, thus each memory cell has an associated byte address.

The byte addresses of the system memory are ordered sequentially from zero through one less than the number of bytes in the memory. Bytes of data are grouped together into half-words, words, double-words, etc. A word is typically defined by the natural data width of the central processing unit (CPU) of the system. If a
20 system has a CPU with a 32 bit wide data bus, for example, then a word may be defined to be 32 bits, arranged as 4 bytes with consecutive addresses. In this example, a half-word is a 16 bit quantity arranged as 2 bytes with consecutive addresses and a double-word is an 64 bit quantity arranged as 8 bytes with consecutive addresses.

Figure 1 shows the layout of a typical system memory. Each rectangle represents a byte of memory. The address for each byte is shown. The figure shows a byte, half-word, word and double-word represented in
25 the memory. These quantities are shown for a system in which a word is defined to be 4 bytes.

Programming languages introduce the notion of data elements as an abstraction of memory cells. That is, programming languages allow a means of associating bytes of storage with data elements. Data elements have associated attributes. One attribute of a data element is the data type, or type, of the data element. A defining characteristic of the data type is the size of the data type. One or more bytes of memory are allocated to
30 contain the values of a given data element depending upon the data type of the data element. The contents, or values, of a data element may be assigned (written) or referenced (read) by using programming language statements. Typical primitive data elements include characters, integers, short integers, long integers, floating-point numbers, strings, pointers, labels, etc.

Additionally, programming languages commonly provide a means for user-definable data elements
35 called data types or data structures. These data structures are defined as a collection of fields, where the fields comprise primitive data elements and other data structures (including recursive references to the data structure itself). Data structures are also commonly used as abstractions of the other devices, such as those mentioned previously, which occupy portions of the system address space. For example, a video display adapter coupled to the computer system may be abstracted as a data structure comprising combination of fields which represent

registers and/or shared memory used to communicate with the adapter, where these hardware elements reside in the address space of and are thus accessible by the CPU.

Data elements in programs have associated memory addresses, that is the byte address of the first byte of the data element. The CPU and other devices in the system supply addresses of the data elements on the address bus in order to read/write the values of the data elements from/to the system memory via the data bus.

As previously discussed, data elements have an associated size, i.e., number of bytes allocated for the data element, based upon the data type of the data element. For example the C language provides an integer data type. In C, an integer data element is represented as the natural word size of the CPU. Hence, an integer on a 32 bit processor is four bytes. Thus, the word shown in Figure 1 could be an integer. The bytes of a four byte word are commonly referred to as the most significant byte (MSB), middle most significant byte (MMSB), middle least significant byte (MLSB) and least significant byte (LSB).

Those who have designed computer systems have been posed with a decision concerning how to take a data element which consists of multiple bytes and order the constituent bytes in memory, that is, how to associate byte addresses with the bytes of a multiple-byte data element. In particular, designers have had to decide the order of the significant bytes of the data element in memory relative to addresses, the order commonly being referred to as byte ordering. Two prevalent formats have been chosen. These two formats are commonly referred to as "little-endian" and "big-endian" format.

Referring now to Figure 2, a four byte integer is shown in both big-endian and little-endian format. In big-endian format the MSB is stored in the lowest of the memory addresses, the MMSB in the next greater memory address, the MLSB in the next greater memory address and the LSB in the next greater memory address. Conversely, in little-endian format the LSB is stored in the lowest of the memory addresses, the MLSB in the next greater memory address, the MMSB in the next greater memory address and the MSB in the next greater memory address.

Referring now to Figure 3, a C language integer, x, is declared and initialized to have the hexadecimal value 0x12345678 and a short, y, is declared and initialized to have the hexadecimal value 0x1234. In this example, an integer is a four byte data element and a short is a two byte data element. The integer and short are shown stored in memory as both a big-endian and little-endian number.

Referring now to Figure 4, a C language string, s, is declared and initialized to be the string "hello". In the C language, strings are conventionally null-terminated. The string is shown stored in memory. The storage of strings is the same for both big-endian and little-endian systems since each character of the string is a single byte, hence there are no byte ordering issues.

Data in a typical computer system is stored in other devices in addition to system memory. Examples of storage devices are permanent storage devices such as disk and tape drives. Additionally, data is stored in registers which exist on devices such as peripheral device controllers. Examples of peripheral device controllers are disk drive controllers, parallel ports, video display adapters, etc. Often these devices need to access data as well as the CPU. Further, the CPU is in communication with these devices and reads and writes data from and to these devices. Still further, the components may be connected by buses which transfer data. These buses also have an associated byte ordering. As long as the individual components of the system which process the data or transfer the data have the same byte ordering, no problems arise. However, if not all the components have the same byte ordering, then data formatting issues arise.

Let us consider an example in which the system CPU is a big-endian processor which submits commands to and receives status from a disk controller via shared data structures stored in system memory. The disk controller is a little-endian controller in that it is coupled to the system via a little-endian expansion bus, such as the Peripheral Component Interconnect (PCI) bus. The embedded processor of the disk controller is a little-endian processor. The CPU places command data structures in the system memory for the disk controller to retrieve and process. The disk controller requires that the command data structures which it retrieves from system memory be stored by the CPU in little-endian format in the system memory. Likewise, the disk controller places status data structures in system memory in little-endian format for the CPU to retrieve and process. The CPU must examine the status data structures with the knowledge that the status data structures are in little-endian format.

The traditional approach to solving problems of the nature identified by the example above has been for the programmer of the software involved, in this case the disk device driver, to painstakingly craft the device driver code to convert the data between different endian formats so as to have the correct byte ordering. This conversion is commonly referred to as "byte swapping", that is, the reversing of the order of the significant bytes of the data.

Typically, the programmer implementing code to perform byte swapping must include additional instructions, namely byte exchanges, shifts, rotates, masks, etc. when loading/storing data elements into/from CPU registers to achieve the byte swapping. This additional burden introduces at least three problems. First, the additional instructions are detrimental to the performance of the system. Second, the code is more bug prone and thus development time is increased. Third, the code developed is less portable.

Therefore, a system and method for converting data between little-endian and big-endian formats is desired which allows software developers to develop more efficient, portable, and bug-free code with respect to byte ordering issues.

Summary of the Invention

The present invention comprises a byte swapping device comprising first and second data ports and data path logic coupled between the first and second data ports. The byte swapping device is employed in a data processing system comprising a data storage device configured to store bytes of data, a processor which reads
5 data from the data storage device and writes data to the data storage device, and the byte swapping device coupled between the data storage device and the processor. The first data port is coupled to the data storage device and the second data port is coupled to the processor. The storage device is typically a system memory or peripheral device controller. The processor processes data in a first endian format, i.e., big-endian or little-endian format, and at least a portion of the data stored in the data storage device is in the opposite byte ordering.

10 The first data port receives first data in a first endian format from the data storage device in response to the processor reading the first data from the data storage device and provides the first data to the data path logic which receives the first data from the first data port and selectively converts the first data from the first endian format to a second endian format and provides the first data to the second data port. The second data port receives the first data from the data path logic and outputs the first data to the processor which receives the first
15 data from the second data port.

Conversely, The second data port receives second data in the second endian format from the processor in response to the processor writing the second data to the data storage device and provides the second data to the data path logic which receives the second data from the second data port and selectively converts the second data from the second endian format to the first endian format and provides the second data to the first data port.

20 The first data port receives the second data from the data path logic and outputs the second data to the data storage device, wherein the data storage device receives the second data from the first data port.

In the preferred embodiment of the invention the address space of the processor is logically divided into data apertures or ranges and the processor provides address signals to the byte swapping device. The byte swapping device selectively converts the first and second data in response to the address signals from the
25 processor based upon whether or not the address is in the data conversion apertures. In one embodiment, the byte swapping device further comprises one or more aperture storage elements configured to store the respective data conversion apertures. The processor programs the data conversion apertures into the aperture storage elements. In another embodiment the data conversion apertures are fixed, that is, not programmable by the processor.

30 In an alternate embodiment, the processor provides control signals to the byte swapping device, wherein the byte swapping device selectively converts the first and second data in response to the control signals from the processor. In one embodiment the processor is configured to execute a characteristic instruction set, wherein the processor provides the one or more control signals to the byte swapping device in response to which instruction in the instruction set the processor executes.

35 Thus, the present invention provides a method for software developers to develop code to execute on a system which performs byte swapping of data between little-endian and big-endian formats thus fostering the development of more efficient, portable, and bug-free code without undue entanglement with byte ordering issues.

Brief Description of the Drawings

A better understanding of the present invention can be obtained when the following detailed description of the preferred embodiment is considered in conjunction with the following drawings, in which:

Figure 1 shows the layout of a typical system memory and common memory entities;

5 Figure 2 shows the layout of significant bytes of a 4 byte word in both big-endian and little-endian format.

Figure 3 shows the storage of a C language integer and short integer in a system memory in both big-endian format and little-endian format;

Figure 4 shows the storage of a C language string in a system memory;

10 Figure 5 shows components of a computer system employing a byte swapping device according to the preferred embodiment of the present invention;

Figure 6 shows components of a computer system employing a byte swapping device according to an alternate embodiment of the present invention;

Figure 7 is a block diagram of the byte swapping device of Figure 5;

15 Figure 8 is a block diagram of the byte swapping device of Figure 6.

Detailed Description of the Preferred Embodiment

Referring now to Figure 5, a computer system 500 employing a byte swapping device according to the preferred embodiment of the present invention is shown. The system 500 comprises a big-endian format processor 510, a storage device 520 and a byte swapping device 530. The processor 510 is coupled to the byte swapping device 530 by a processor data bus 540. The storage device 520 is coupled to the byte swapping device 530 by a system data bus 550. The processor 510, storage device 520 and byte swapping device 530 are coupled to a processor address bus 560. The processor 510, the storage device 520 and the byte swapping device 530 are coupled to a control bus 570.

The storage device 520 is representative of a system memory, peripheral device, peripheral device controller, or other elements capable of storing multiple bytes of data. At least a portion of the data stored in the storage device 520 comprises data elements stored in little-endian format. Other storage devices (not shown) may be coupled to the system data bus 550 and the processor address bus 560. Signals representative of common computer system control bus signals are transmitted on the control bus 570. The control bus signals comprise a clock signal, a read/write signal, memory/IO signal, a code/data signal, byte enable signals, ready signal, data width signals, and an address strobe signal.

In the system 500 shown, the processor data bus 540 and system data bus 550 are 32 bits wide organized in four sets of eight bits, commonly referred to as byte lanes. However, it is noted that data buses of different widths are contemplated in the present invention and the invention is intended to extend to data buses of larger or smaller numbers of bits. The processor 510 reads data from the storage device 520 and writes data to the storage device 520.

As the processor 510 reads data from the storage device 520 the data passes from the storage device 520 along the system data bus 550 to the byte swapping device 530. The byte swapping device 530 receives the data from the system data bus 550 and selectively byte swaps the data. The byte swapping device 530 then provides the data to the processor 510 on the processor data bus 540. The processor 510 receives the data from the processor data bus 540 into registers within the processor 510 (not shown).

Conversely, as the processor 510 writes data to the storage device 520 the data passes from the processor 510 along the processor data bus 540 to the byte swapping device 530. The byte swapping device 530 receives the data from the processor data bus 540 and selectively byte swaps the data. The byte swapping device 530 then provides the data to the storage device 520 on the system data bus 550. The storage device 520 receives the data from the system data bus 550 into storage locations within the storage device 520 (not shown).

The processor 510 provides the address of the data to be read from or written to the storage device 520 on the processor address bus 560. The processor 510 also provides the necessary control signals on the control bus 570 to instruct the storage device 520 as to how many bytes of data are to be transferred, the direction of the transfer, on which byte lanes the data is to be transferred, etc. The storage device 520 uses the address supplied on the processor address bus 560 and the control signals supplied on the control bus 570 to determine which data to provide or receive. At least some of the data stored in the storage device 520 is little-endian format data and is byte swapped into big-endian format by the byte swapping device 530 to be provided to the processor 510 when the processor 510 reads the data from the storage device 520. Conversely, at least some of the data written by the processor 510 is in big-endian format and is byte swapped by the byte swapping device 530 into little-

endian format to be provided to the storage device 520 when the processor 510 writes the data to the storage device 520.

5 In the preferred embodiment of the present invention the byte swapping device 530 selectively performs byte swapping of the data transferred between the processor 510 and the storage device 520 based upon the address signals provided on the processor address bus 560 and the control signals provided on the control bus 570 by the processor 510.

The byte swapping device 530 receives the address signals from the processor address bus 560 and selectively performs byte swapping based upon whether or not the address on the address signals falls within one or more data conversion apertures, i.e., ranges of contiguous byte addresses on the processor address bus 560. 10 Thus the byte swapping device 530 advantageously provides a hardware mechanism for performing byte swapping in the system 500.

In one embodiment, the one or more data conversion apertures are fixed within the byte swapping device 530. In one embodiment in which the data conversion apertures are fixed, the byte swapping device 530 examines the most significant bits of the address to determine whether or not to swap the data bytes. In another 15 embodiment, the one or more data conversion apertures are programmable by the processor 510.

Alternate embodiment

Referring now to Figure 6, a computer system 500 employing a byte swapping device according to an alternate embodiment of the present invention is shown. The system 500 is similar to the system 500 of Figure 5 20 and corresponding elements are numbered identically for clarity. The alternate embodiment further comprises byte swap control signals 680 provided by the processor 510 to the byte swapping device 530. The byte swapping device 530 selectively performs byte swapping of the data transferred between the processor 510 and the storage device 520 based upon the byte swap control signals 680 provided to the byte swapping device 530 by the processor 510.

25 The processor 510 is configured to execute a characteristic instruction set. The instruction set includes instructions to perform transfers of data between the processor 510 and other elements of the system 500, such as the storage device 520 and to perform byte swapping on the data being transferred. The processor generates the byte swap control signals 680 based upon the given instruction being executed by the processor 510. If the processor 510 executes a byte swapping data transfer instruction the processor 510 generates control signals 680 30 to instruct the byte swapping device 530 to swap the data bytes as they pass through the byte swapping device 530. If the processor 510 executes a non-byte swapping instruction the processor 510 generates control signals 680 to instruct the byte swapping device 530 not to swap the data bytes as they pass through the byte swapping device 530.

It is noted that although embodiments are described which operate on one, two and four byte data types, 35 the invention contemplates embodiments which operate on data types of other sizes, in particular larger data types such as eight byte data types which are common in currently existing processors.

It is further noted that although embodiments are described in which the data processing device is a big-endian format device and the data storage device is a little-endian format device, an embodiment is contemplated in which the two devices have reversed endian formats.

It is further noted that although the embodiments are described with a single processor, alternate embodiments are contemplated which comprise multiple processors.

It is further noted that although embodiments are described in reference to computer systems, the invention contemplates embodiments in other digital systems which store data in both big-endian and little-endian formats.

The Byte Swapping Device

Referring now to Figure 7, a block diagram of the byte swapping device 530 of Figure 5 according to the preferred embodiment of the present invention is shown. The byte swapping device 530 comprises a first data port 710 coupled to the storage device 520 (of Figure 5) and a second data port 720 coupled to the processor 510 (of Figure 5). Data path logic 730 is coupled between the first data port 710 and second data port 720.

The data ports comprise four byte-wide bi-directional bus transceivers. The data path logic 730 comprises two sets of 4 byte to 1 byte multiplexers. Each multiplexer output of the first set of multiplexers is coupled to a respective transceiver of the second data port 710. Each multiplexer output of the second set of multiplexers is coupled to a respective transceiver of the first data port 710. The four inputs of each of the first set of multiplexers is coupled to a respective bus transceiver of the second data port 720. The four inputs of each of the second set of multiplexers is coupled to a respective bus transceiver of the first data port 710. The multiplexer select signals are provided by multiplexer select logic 740. The transceivers receive the read/write signal (not shown) from the control bus 580 to determine the direction of data flow through the transceivers.

The first data port 710 receives data from and provides data to the storage device 520. The second data port 720 receives data from and provides data to the processor 510. Each of the transceivers of the first data port 710 provides data to the second set of multiplexers of the data path logic 730. Each of the transceivers of the second data port 720 provides data to the first set of multiplexers of the data path logic 730. The multiplexers select one of the four bytes of data received based upon the multiplexer select signals provided by the multiplexer select logic 740. Thus the byte swapping device 530 advantageously provides a hardware mechanism for performing byte swapping.

The byte swapping device 530 of Figure 7 is configured to perform byte swapping on half-words (i.e., two byte data entities) and words (i.e., four byte data entities). Other embodiments are contemplated which are configured to perform byte swapping on eight, sixteen and larger power of two byte entities.

In the preferred embodiment of the byte swapping device 530 the multiplexer select logic 740 receives the address signals from the processor address bus 560 and uses the address signals to generate the multiplexer select signals. The byte swapping device 530 selectively performs byte swapping on the data passing through the byte swapping device 530 by comparing the address on the address signals with the one or more data conversion apertures as described in the description of Figure 5.

Alternatively, the byte swapping device 530 further comprises a plurality of aperture storage elements (not shown) which store the one or more data conversion aperture values. The aperture storage elements are coupled to the multiplexer select logic 740 and provide the multiplexer select logic 740 with the one or more data conversion aperture values. The processor 510 programs the aperture storage elements with the desired data conversion aperture values.

Referring now to Figure 8, a block diagram of the byte swapping device 530 of Figure 6 according to the alternate embodiment of the present invention is shown. The byte swapping device 530 is similar to the byte swapping device 530 of Figure 7, except that the multiplexer select logic 740 receives the byte swap control signals 680 and uses the byte swap control signals 680 to generate the multiplexer select signals.

5

Source Code Translator

The present invention contemplates a source code translator, or compiler, which translates source code modules written in a high level language, such as the C language, Pascal, Ada, or other languages, into one or more object modules. The object modules are then linked together into an executable computer program to be executed on the computer system 500. The compiler is similar to high-level language compilers known in the art of compilers. However, the compiler of the present invention further comprises a feature for generating code which advantageously employs the byte swapping capability of the computer system 500. The compiler is cognizant of the data conversion apertures of the computer system 500.

The compiler receives the string of characters comprising the source code module and parses the characters into tokens. At least a portion of the tokens are execution statements defined by the high-level language and identifiers associated with data elements. The compiler generates instructions chosen from the instruction set of the processor 510 to perform the execution statements.

The compiler also generates an object offset for each of the data elements so that the instructions may correctly access the data elements. The value of the object offset for a given data element depends upon several factors. One factor is whether the compiler is generating relocatable code or absolute code. Another factor is whether or not the compiler is generating code for a processor which uses base registers. Another factor is the nature of the data element, particularly whether static or dynamic storage allocation is being performed. For example, if the data element is allocated on stack, the object offset will be relative to a stack pointer. The generation of object offsets for data elements is well known in the art of compilers.

The compiler also advantageously generates a format base for each data element. The compiler examines the data type of each data element and decides whether or not transfers of the data element between the processor 510 and the storage device 520 require byte swapping by the byte swapping device 530. If the data element requires byte swapping, such as an integer, short integer, long integer, or floating-point number, the compiler generates a format base to effect byte swapping. If the data element does not require byte swapping, such as a single character within a string, the compiler generates a format base to not effect byte swapping. The format base is added to the object offset to calculate a data aperture offset for each data element as shown below.

$$\text{data aperture offset} = \text{object offset} + \text{format base}$$

A typical compiler stores the object offset associated with each data element in the object code module so that a linker may use the object offset to generate executable code and a loader may load the executable code into a system memory for execution on the computer system 500. However, the compiler of the present invention stores the data aperture offset in the object code module instead of the object offset for each data element. Hence, the object module advantageously contains address offsets for data elements which are aware of the data conversion apertures of the system 500 previously described.

40

The object module is linked into an executable program. The executable program is loaded into the computer system 500 and executed. Each data element has a base memory address bound by the system in the storage device 520, typically in system memory or registers or shared memory of a peripheral device. The binding of the base memory address to the data element is achieved in one of at least three ways.

5 In the first scenario, the storage space for the data element is dynamically allocated to the executable program by an operating system running on the computer system at the request of the program. In this case the operating system keeps track of free areas of system memory and allocates a portion of the free system memory to the program to store the data element. The operating system returns the base memory address of the allocated memory to the program, which the program uses to access the data element. Hence, the base memory address is
10 bound to the data element during the execution, or run-time, of the program.

In the second scenario, the storage space for the data element is allocated to the data element by the loader when the program is loaded. The base address of the storage space is communicated to the program by the loader. Hence, the base memory address is bound to the data element during the loading, or load-time, of the program.

15 In the third scenario, the program, and thus data elements, are loaded at a fixed locations in the system memory as determined by the linker when creating the executable program. Hence, the base memory address is bound to the data element during the linking, or compile-time, of the program.

The format base selectively effects byte swapping as follows. As the program executes the processor 510 executes instructions which perform transfers of the data elements between the processor 510 and the
20 storage device 520. The processor 510 computes the addresses of the data elements, referred to in this disclosure as data aperture addresses, for each element by adding the base memory address for the data element and the data aperture offset as shown in the following equation.

$$\text{data aperture address} = \text{base memory address} + \text{data aperture offset}$$

25 The processor 510 generates the data aperture address of a given data element on the processor address bus 560 to read or write the data element. The byte swapping device 530 receives the data aperture address from the address bus 560 and selectively byte swaps the data element as the data element is read from the storage device 520 to the processor 510 or written from the processor 510 to the storage device 520 based upon the
30 relationship between the value of the data aperture address received and the data conversion apertures defined for the byte swapping device 530 as discussed in relation to Figure 5.

In the preferred embodiment, the format bases are chosen such that different format bases are distinguished by using the most significant bits of the address. The number of significant bits used is the ceiling function of the base two logarithm of the number of required data conversion apertures as shown in the following
35 equation.

$$\text{Number of address bits} = \text{ceiling}[\log_2(\text{number of data conversion apertures})]$$

Hence, an implementation requiring 5 data conversion apertures requires 3 bits of address.

5 In one embodiment of the compiler, the compiler generates a list of data conversion aperture values which are stored in the object module and subsequent executable program. The loader executing on the processor 510 programs the data conversion apertures in the aperture storage elements of the byte swapping device 530 before executing the program.

Conclusion

10 Therefore, the present invention comprises a byte swapping device employed within a computer system for automatically converting data between big-endian and little-endian formats for the various elements of the computer system. The invention comprises a compiler which advantageously generates object code modules to be linked into executable programs for execution on the computer system employing the byte swapping apparatus.

15 Although the system, method and apparatus of the present invention have been described in connection with the preferred embodiment, it is not intended to be limited to the specific form set forth herein, but on the contrary, it is intended to cover such alternatives, modifications, and equivalents, as can be reasonably included within the spirit and scope of the invention as defined by the appended claims.

Claims

1. A byte swapping device, comprising:

first and second data ports; and

data path logic coupled between said first and second data ports;

5 wherein said first data port receives first data in a first endian format and provides said first data to said data path logic;

wherein said data path logic receives said first data from said first data port and selectively converts said first data from said first endian format to a second endian format and provides said first data to said second data port;

10 wherein said second data port receives said first data from said data path logic and outputs said first data;

wherein said second data port receives second data in said second endian format and provides said second data to said data path logic;

15 wherein said data path logic receives said second data from said second data port and selectively converts said second data from said second endian format to said first endian format and provides said second data to said first data port;

wherein said first data port receives said second data and outputs said second data;

wherein said first and second endian formats have different byte ordering.

20 2. The byte swapping device of claim 1 further comprising one or more control signal inputs;

wherein said data path logic selectively converts said first and second data in response to said one or more control signal inputs.

3. The byte swapping device of claim 1 further comprising address signal inputs;

25 wherein said data path logic selectively converts said first and second data in response to said address signal inputs.

4. The byte swapping device of claim 3, wherein said data path logic converts said first and second data only when said address signal inputs have an address in one or more data conversion apertures.

30 5. The byte swapping device of claim 4 further comprising one or more aperture storage elements configured to store respective said one or more data conversion apertures.

6. The byte swapping device of claim 5, wherein a data processing device coupled to said byte swapping device writes said one or more data conversion apertures to said one or more aperture storage elements.

7. A data processing system, comprising:

a data storage device configured to store bytes of data;

a data processing device which reads data from said data storage device and writes data to said data storage device; and

a byte swapping device comprising data path logic coupled between first and second data ports, wherein said first data port is coupled to said data storage device, wherein said second data port is coupled to said data processing device;

wherein said first data port receives first data in a first endian format from said data storage device in response to said data processing device reading said first data from said data storage device and provides said first data to said data path logic;

wherein said data path logic receives said first data from said first data port and selectively converts said first data from said first endian format to a second endian format and provides said first data to said second data port;

wherein said second data port receives said first data from said data path logic and outputs said first data to said data processing device, wherein said data processing device receives said first data from said second data port;

wherein said second data port receives second data in said second endian format from said data processing device in response to said data processing device writing said second data to said data storage device and provides said second data to said data path logic;

wherein said data path logic receives said second data from said second data port and selectively converts said second data from said second endian format to said first endian format and provides said second data to said first data port;

wherein said first data port receives said second data from said data path logic and outputs said second data to said data storage device, wherein said data storage device receives said second data from said first data port;

wherein said first and second endian formats have different byte ordering.

8. The system of claim 7, wherein said data processing device provides one or more control signals to said byte swapping device, wherein said byte swapping device selectively converts said first and second data in response to said one or more control signals from said data processing device.

9. The system of claim 8, wherein said data processing device is configured to execute a characteristic instruction set, wherein said data processing device provides said one or more control signals to said byte swapping device in response to which instruction in said instruction set said data processing device executes.

10. The system of claim 7, wherein said data processing device provides address signals to said byte swapping device, wherein said byte swapping device selectively converts said first and second data in response to said address signals from said data processing device.

11. The system of claim 10, wherein said byte swapping device converts said first and second data only when said address signals provide an address in one or more data conversion apertures.

12. The system of claim 11, wherein said byte swapping device further comprises one or more aperture storage elements configured to store respective said one or more data conversion apertures.

5 13. The system of claim 12, wherein said data processing device writes said one or more data conversion apertures to said one or more aperture storage elements.

10 14. A method for performing byte swapping on data in a data processing system, wherein said data processing system comprises a data storage device configured to store bytes of data, a data processing device which reads data from said data storage device and writes data to said data storage device, and a byte swapping device comprising data path logic coupled between first and second data ports, wherein said first data port is coupled to said data storage device, wherein said second data port is coupled to said data processing device, comprising:

15 said first data port receiving first data in a first endian format from said data storage device in response to said data processing device reading said first data from said data storage device;

said first data port providing said first data to said data path logic;

said data path logic receiving said first data from said first data port;

said data path logic selectively converting said first data from said first endian format to a second endian format;

20 said data path logic providing said first data to said second data port;

said second data port receiving said first data from said data path logic;

said second data port outputting said first data to said data processing device;

said data processing device receiving said first data from said second data port;

25 said second data port receiving second data in said second endian format from said data processing device in response to said data processing device writing said second data to said data storage device;

said second data port providing said second data to said data path logic;

said data path logic receiving said second data from said second data port;

said data path logic selectively converting said second data from said second endian format to said first endian format;

30 said data path logic providing said second data to said first data port;

said first data port receiving said second data from said data path logic;

said first data port outputting said second data to said data storage device,

said data storage device receiving said second data from said first data port;

wherein said first and second endian formats have different byte ordering.

35 15. The method of claim 14, wherein said byte swapping device selectively converting said first and second data is in response to one or more control signals provided to said byte swapping device by said data processing device.

16. The method of claim 15, wherein said data processing device is configured to execute a characteristic instruction set, wherein said data processing device provides said one or more control signals to said byte swapping device in response to which instruction in said instruction set said data processing device executes.

5

17. The method of claim 14, wherein said data path logic selectively converting said first and second data is in response to address signals received by said byte swapping device from said data processing device.

10

18. The method of claim 17, wherein said data path logic converting said first and second data occurs when said address signals have an address in one or more data conversion apertures.

19. The method of claim 18, wherein said byte swapping device further comprises one or more aperture storage elements configured to store respective said one or more data conversion apertures.

15

20. The method of claim 19, wherein said data processing device writes said one or more data conversion apertures to said one or more aperture storage elements.




Fig. 1

2/8

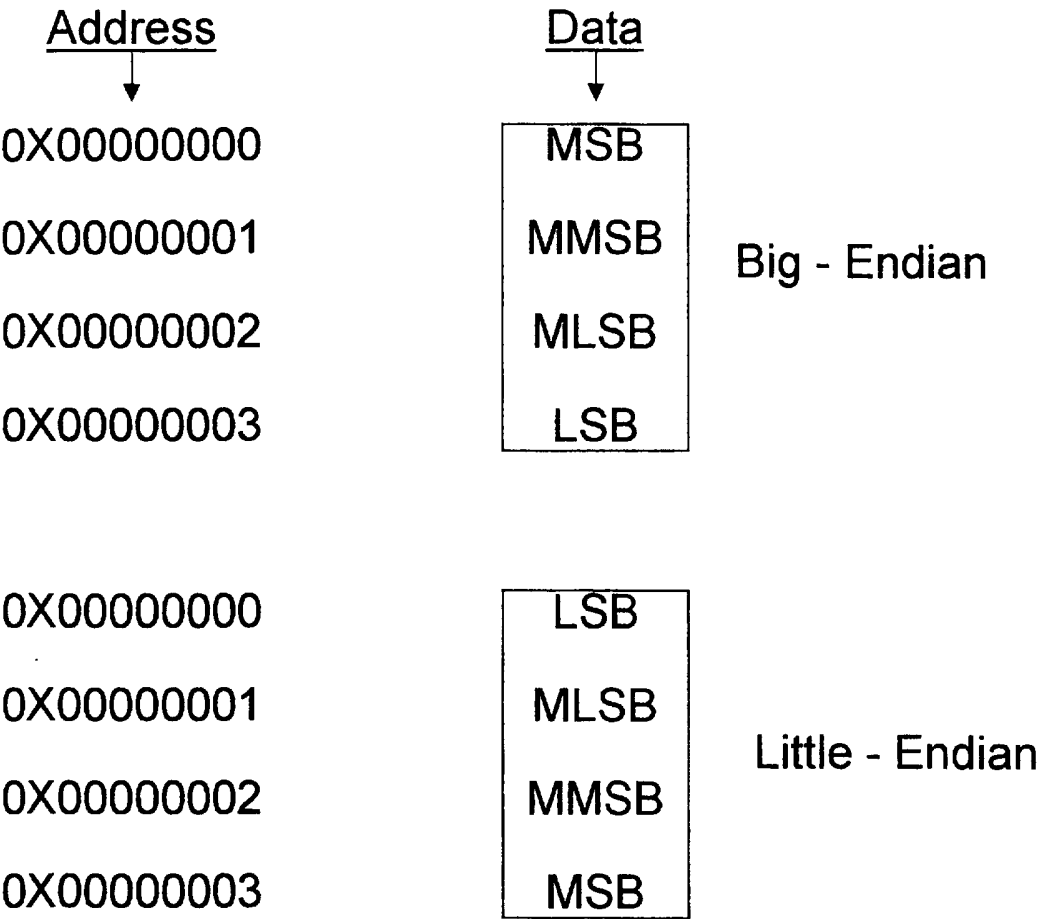


Fig. 2

3/8

```
int x = 0X12345678;
```

<u>Address</u>	<u>Data</u>	
0X00000000	0X12	Big - Endian
0X00000001	0X34	
0X00000002	0X56	
0X00000003	0X78	

0X00000000	0X78	Little - Endian
0X00000001	0X56	
0X00000002	0X34	
0X00000003	0X12	

```
short y = 0X1234
```

0X00000000	0X12	Big - Endian
0X00000001	0X34	

0X00000000	0X34	Little - Endian
0X00000001	0X12	

Fig. 3

4/8

```
char s[6] = "hello";
```

<u>Address</u>	<u>Data</u>
0X00000000	'h'
0X00000001	'e'
0X00000002	'l'
0X00000003	'l'
0X00000004	'o'
0X00000005	Ø

Fig. 4

5/8




Fig. 5

6/8


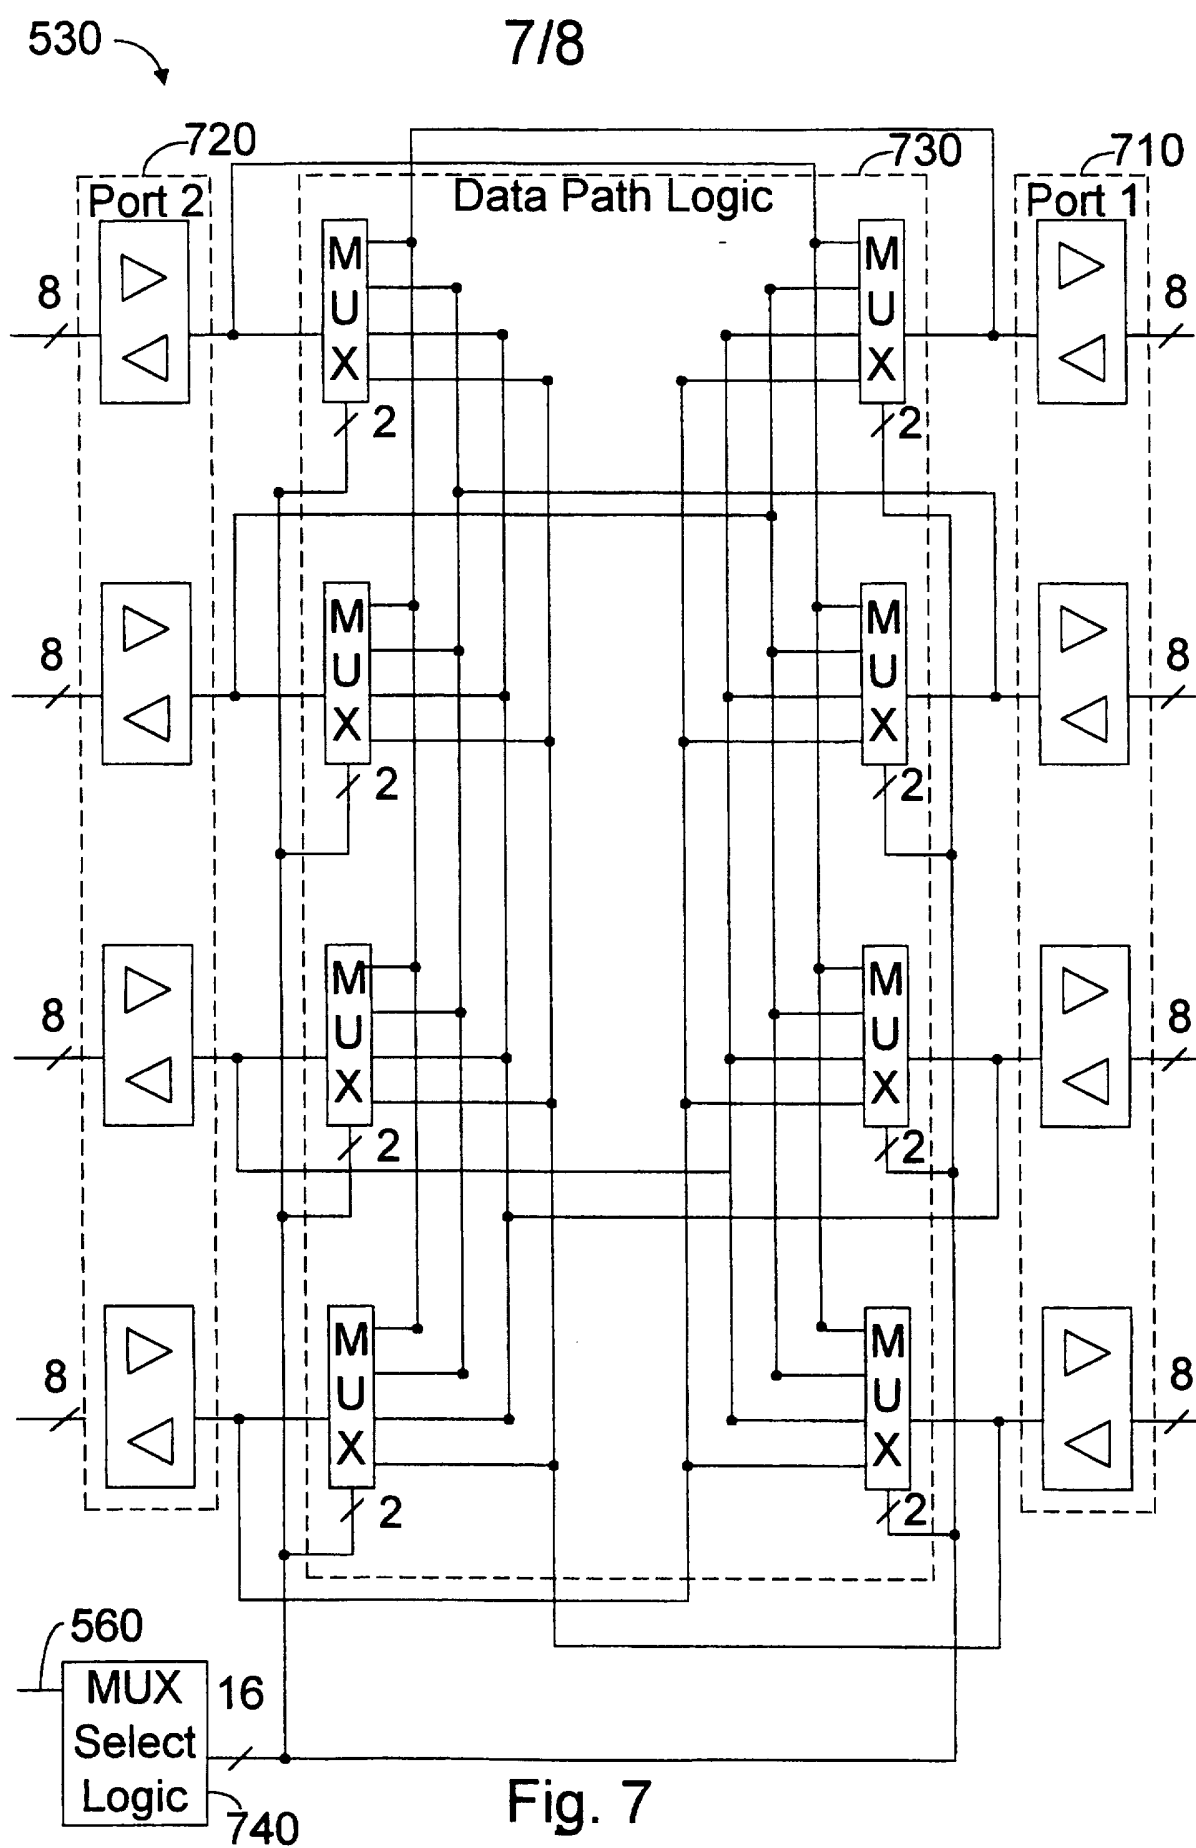
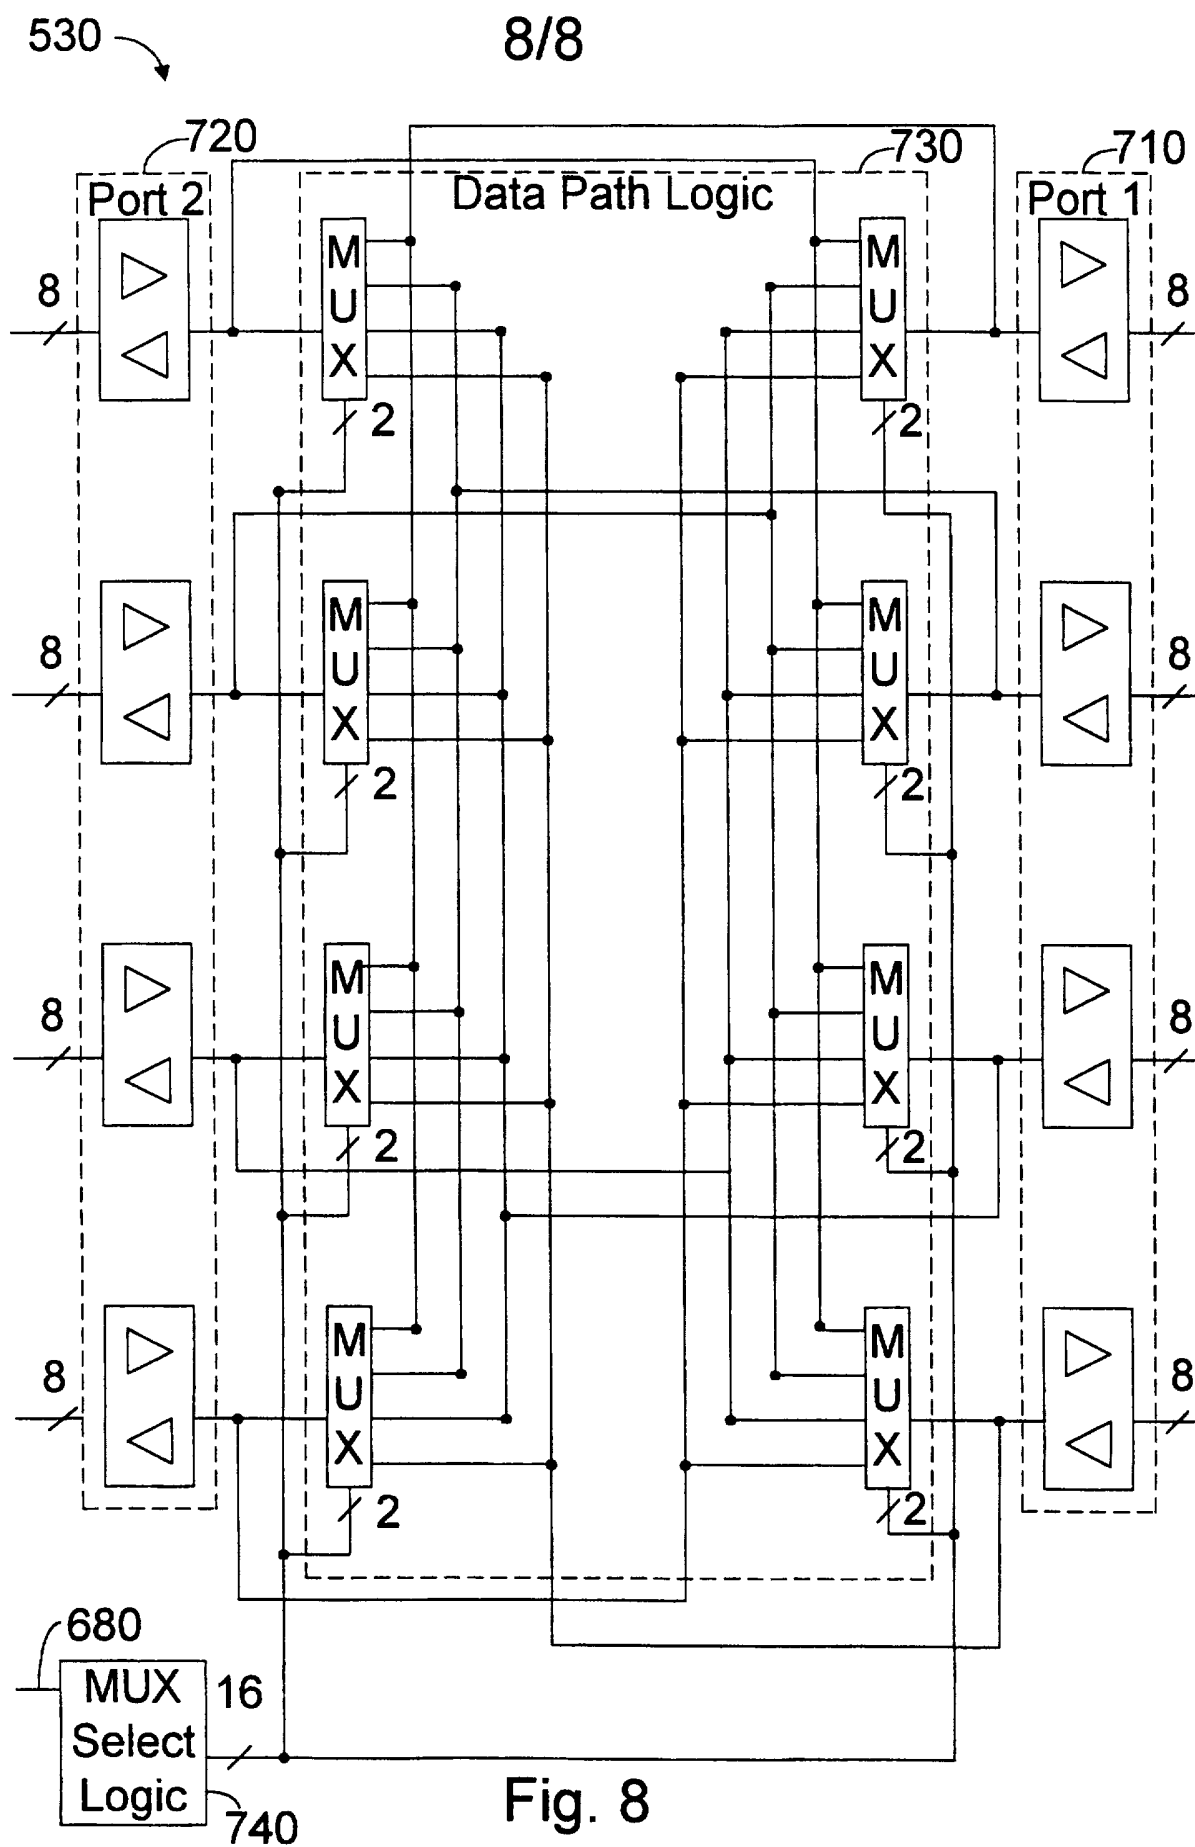


Fig. 6



8/8



INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 97/00891

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F13/40

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 94 15269 A (OLIVETTI) 7 July 1994	1,2,7-9, 14-16
Y	see page 16, line 31 - page 31, line 4; figures 7-16	3-6, 10-13, 17-20
Y	IBM TECHNICAL DISCLOSURE BULLETIN, vol. 39, no. 1, January 1996, NEW YORK US, pages 255-256, XP000556393 "Using Multiple Preallocated Pools of Memory for Bi-Endian Systems" see the whole document	3-6, 10-13, 17-20



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

27 May 1997

Date of mailing of the international search report

05.06.97

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+ 31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+ 31-70) 340-3016

Authorized officer

Gill, S

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No.

PC I/US 97/00891

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9415269 A	07-07-94	EP 0629303 A	21-12-94
		JP 7505972 T	29-06-95
