

AUDIO EFFECTS PROCESSOR HAVING DECOUPLED INSTRUCTION EXECUTION AND AUDIO DATA SEQUENCING

BACKGROUND OF THE INVENTION

The present invention relates to an audio signal processor. More specifically, the present invention relates to an audio effects processor integrated on a single chip and including memory and architecture capable of implementing audio delay lines independent of the operation of an instruction execution unit of the effects processor.

Sound effects audio signal processors (ASPs) are commonly used to generate sound effects in a multitude of audio components, some examples being programmable musical instruments, video games, cinema sound systems, virtual reality systems, and computer audio systems. Such sound effects include, but are not limited to, reverberation effects, 3-D audio, and distortion effects. ASPs create sound effects by executing programs containing a series of instructions. Each instruction directs the ASP to perform a specific logical and/or arithmetic operation on the received signal, resulting in the program's creation of a particular sound effect.

It is common for ASPs to implement sound effects through the use of delay lines, which often take the form of a large block of relatively slow "tank" RAM (TRAM) memory. Some known ASPs perform operations on the TRAM data through a TRAM engine under the control of an instruction execution unit of the ASP. That is, when the instruction execution unit wants to operate on the signal output from a delay line, it schedules a TRAM read operation to begin at a specific instruction cycle, waits for a specific number of cycles for the read operation to occur and then performs an operation to retrieve the data.

In processors of this type, the instruction execution unit must be tightly coupled to the TRAM engine because any mis-synchronization in the scheduling of operations could result in the loss of the TRAM data. Also, valuable processor time of the instruction execution unit is used in coordinating and implementing the sound effects. Accordingly, improvements to the architecture of such ASPs are desirable.

SUMMARY OF THE INVENTION

The present invention provides an audio effects processor in which the TRAM engine (hereinafter referred to as the "sound memory engine") is decoupled from the instruction execution unit of the processor. In the present invention, the sound memory engine and instruction execution unit operate independent of each other, passing TRAM data through a shared area of a memory. The sound memory engine reads and writes TRAM data according to its own schedule, using the shared memory area as a source for TRAM writes and a destination for TRAM reads. When the executing program requires TRAM data, the instruction execution unit reads or writes to these same shared register file locations.

A sound effects processor according to the present invention includes a first memory, an instruction execution unit and a sound memory engine integrated onto a single chip. The first memory includes a first address space that is addressable by both the instruction execution unit and the sound memory engine. The instruction execution unit executes sound processing instructions, including instructions to read from and write to the first address space, and the sound memory engine operates independent of the program counter (instruction execution rate) of the instruction execution unit to read from and write to the first address space and perform address arithmetic and memory accesses to a second memory as necessary to implement audio data storage (e.g., delay lines whose delay length can be varied dynamically -- a capability required for many audio effects like chorusing, flanging and pitch shifting). The second memory can be internal or external to the effects processor.

For a further description of the nature and advantages of the present invention, reference should be made to the following description taken in conjunction with the accompanying drawings.

5

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1A depicts a representative multimedia personal computer system in which the audio effects processor of the present invention may be employed;

Fig. 1B depicts a simplified representation of the internal architecture of the multimedia personal computer system depicted in Fig. 1A;

10

Fig. 2 is a simplified block diagram of multimedia board 28 shown in Fig. 1B, a board onto which the audio signal processor of the present invention could be incorporated;

Fig. 3 is a simplified block diagram of one embodiment of the audio signal processor shown in Fig. 2;

15

Fig. 4 is a simplified block diagram of audio effects processor 50 in accordance with the present invention shown in Fig. 3;

Fig. 5 is a block diagram of one embodiment of sound memory engine 74 shown in Fig. 4;

Fig. 6 is a block diagram showing the allocation of memory space within one embodiment of memory 72 shown in Fig. 4;

20

Figs. 7A-7D are diagrams that illustrate the circular addressing scheme used for generating delay lines according to the present invention; and

Figs. 8A-8D are figures illustrating how sound memory engine 74 prevents TRAM read and write hazards that would otherwise be possible because memory 72 is shared with instruction execution unit 70.

25

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The present invention is applicable to digital sound generation systems of all kinds. Advanced audio effects can be provided to video games, multimedia computer systems, virtual reality environments, cinema sound systems, home theater, and home digital audio systems, for example. Fig. 1A depicts a representative multimedia personal computer 10 with a monitor 12 and left and right

30

speakers 14 and 16, an exemplary system that can be enhanced in accordance with the invention with sound effects such as three-dimensional audio.

Fig. 1B depicts a greatly simplified representation of the internal architecture of personal computer 10. Personal computer 10 includes a CPU 18, memory 20, a floppy drive 22, a CD-ROM drive 24, a hard drive 26, and a multimedia card 28. Each of the components of computer 10 shown in Fig. 1B communicate with each other over a bus system 29. Of course, many possible computer configurations could be used with the invention. In fact, the present invention is not limited to the context of personal computers and finds application in video games, cinema sound systems and many other systems.

Fig. 2 illustrates a typical multimedia card 28 on which an integrated circuit according to the present invention may be mounted. Multimedia card 28 includes a sound processor chip 30 mounted on a circuit board 32. As shown in Fig. 2, a CDROM connector 34, an AC97 CODEC 36, an optional AC3 decompressor/decoder 38 and a mixer 39 are all connected to sound processor chip 30 through appropriate interfaces.

Also shown in Fig. 2 are various other connections to sound processor 30, including a joystick connector 42, a phone line connection 44 for a modem (not shown), a line-in connector 46, a microphone connector 47, and a speaker output 48. In addition, a connection to a PCI bus 49, which is part of bus system 29, is shown. Bus 49 connects to the host microprocessor 18 and to main memory 20.

A simplified block diagram of sound processor 30 is shown in Fig. 3. Sound processor 30 includes three primary functional units: a sound processing engine 40, a sound effects engine 50 and a host interface unit 60. Sound processing engine 40 is a 64-voice wavetable synthesizer that employs an eight-point interpolation algorithm for professional quality audio playback as described in U.S. Patent No. 5,342,990 entitled "Digital Sampling Instrument Employing Cache Memory" and 16 summing effects send buses. Each of the 64 voice channels can be routed, at its own programmable amplitude, to an arbitrary selection of four of these buses. Host interface unit 60 interfaces sound processor 30 with host CPU 18 using the PCI protocol. Sound effects engine 50 receives input from sound

processing engine 40 and from additional audio inputs such as CD Audio, ZVideo, a microphone jack, a stereo input and an auxiliary S/PDIF input among others. Further details of sound effects engine 50 and host interface unit 60 are described below with respect to Fig. 4. Other details of host interface unit 60 along with some details of sound processing engine 40 and other portions of sound processor 30 are set forth in U.S. Serial No. ____/____ (attorney docket 17002-86) entitled "AUDIO EFFECTS PROCESSOR WITH MULTIPLE ASYNCHRONOUS AUDIO STREAMS," having David P. Rossum and Scott Fuller as inventors and assigned to Creative Technologies, Ltd., the assignee of the present invention. The 5,342,990 patent and the ____/____ application (attorney docket 17002-86) are both hereby incorporated by reference in their entirety.

Audio Effects Processor 50

According to the present invention, sound effects engine 50 (also referred to as "effects processor 50") includes separate functional units to 1) execute audio signal processing instructions and 2) implement audio data storage (e.g., audio signal delay lines and table look-ups). These functional units operate independent of each other and exchange data through a shared memory address space. One particular implementation of an audio effects processor of the present invention is illustrated in Fig. 4 and discussed below. In no respect is the present invention limited to this specific implementation, however. After reading the following description, a person of ordinary skill in the art will understand that other implementations of effects processor 50 are possible without departing from the concept of the present invention.

Fig. 4. is a simplified block diagram of one embodiment of effects processor 50 shown in Fig. 3 above. The basic architecture of effects processor 50 combines an instruction execution unit 70, a high-speed internal memory 72, and a sound memory engine 74 (also referred to as "TRAM engine 74"), which interfaces to internal and external slower-speed, high capacity TRAM memories 76 and 78. Audio signals are received by effects processor 50 at a processor input 71, where they are directed to a dual buffered portion 83 of memory 72. In this embodiment, the audio signals received at input 71 include 16 voices output from sound

processing engine 40, 2 lines for a stereo input, 2 lines for a CD audio input, 2 lines for a zvideo input, a microphone input, 2 lines for an auxiliary S/PDIF input and a 6-line I²S input.

5 Instruction Execution Within Effects Processor 50

Instruction execution unit 70 executes instructions stored in an internal microprogram memory 80, separate from memory 72 to perform particular operations on one or more of the audio signals stored in memory address space 83. Each instruction has a 4-bit opcode field, three 10-bit input operand address fields
10 and a 10-bit output operand address field. The opcode field holds the instruction's operation code (opcode) that specifies the instruction's operation to the processor. The operand address fields hold addresses to spaces in memory 72. Because the operands are 10-bits long in this embodiment, the maximum addressable space of memory 72 available to DSP instructions is 1024 locations. All 1024 locations of
15 memory 72 are addressable.

Instruction execution unit 70 retrieves data stored in the locations in memory 72 from the locations specified in the instruction's three input operands over input data lines 75, 77 and 79. Instruction execution unit 70 then performs the requested operation specified by the instruction's opcode and writes the resultant
20 data over output data line 81 to the memory address specified in result operand of the received instruction. The resultant data may be directed to an audio output port 87 by writing to one of the 32 output buffers in memory space 83, stored in another portion of memory 72, or reused in subsequent operations by instruction execution unit 70.

25 Instruction execution within instruction execution unit 70 is sample-locked and memory-based, rather than register-based. Sample-locked execution means that instruction unit 70 executes a single pass through microprogram memory 80 every sample period. Thus, instruction execution is inherently tied to the system's audio sample rate. In this embodiment, the sample
30 rate is 48 KHz as derived from AC97 Codec 36 shown in Fig. 2. In other embodiments this sample rate may be a different set frequency or may be a variable

frequency. For example, in another embodiment, the sample rate is in the range of about 40-50 KHz as derived from a system clock on board 32 (not shown).

Memory-based instruction execution means that operands are fetched directly from memory 72 without the bottleneck of an input register set at the input to the instruction execution unit. It also means that there are no volatile output registers that need to be flushed after each instruction in order to save results.

Instruction execution unit 70 executes all the instructions in microprogram memory 80 "in-line" without branches, loops or subroutine calls. Such in-line execution ensures that all processing within the basic real-time deadline (one sample period) is achieved by design. It also simplifies the design of the execution pipeline, since instruction sequencing is completely deterministic, i.e., there are no branches in the flow of control.

Although instruction execution always proceeds in-line from the start of microprogram memory 80 to the end of the memory, there are provisions for the conditional execution of instructions. Details of such provisional instruction execution along with specifics of the instruction set executed by instruction execution unit 70 are set forth in U.S. Serial No. ____/____ (attorney docket 17002-84) entitled, "Processor With Instruction Set for Audio Effects," having Stephen Hoge as the inventor and assigned to Creative Technologies, LTD., the assignee of the present invention. The ____/____ application (attorney docket 17002-84) is hereby incorporated by reference for all purposes.

Each instruction executes in a single instruction cycle whose period is the sample rate divided by the number of instructions in microprogram memory 80. Thus, in an embodiment where microprogram memory 80 is 512 instructions long and the sample rate is 48 KHz, the instruction cycle is 40.7 nanoseconds long. The results of each instruction are available for use in the following instruction cycle.

Host Interface Unit 60

Also shown in Fig. 4 is host interface unit 60, which allows the host processor (CPU 18) to control the operation of audio effects processor 50. Such control is achieved by allowing the host to initialize and read and write data and

executable instructions to memory 72 and/or microprogram memory 80. Such memory read and write operations can be executed transparent to the execution of digital signal processing (DSP) programs allowing audio effects processor 50 to support simultaneous startup, execution and shutdown of multiple independent and separately loaded programs. In this embodiment, communication between the host processor and audio effects processor 50 through interface unit 60 occurs over a PCI bus using a PCI protocol, but other protocols can be implemented in other embodiments. Also in still other embodiments, host interface unit 60 is a direct, memory-mapped microprocessor interface suitable for use with either big-endian or little-endian processors.

Sound Memory Engine 74

Sound memory engine 74 is the interface between memory 72 and the large capacity TRAM memory used for long-term audio data storage. Fig. 5 is a block diagram of one embodiment of sound memory engine 74 configured to implement audio data storage in the form of audio delay lines. As shown in Fig. 5, the primary components of sound memory engine 74 are control bits 83, the TRAM data and address buffers 84 and 85 (part of memory 72), a delay base address counter 86, an align control circuit 87, an address generator 88 and a data formatter 90. The address generator is responsible for circular addressing for generating delay lines. The data formatter is used to compress 32-bit operands in memory 72 into the narrower word widths implemented by the physical TRAM during TRAM write operations and then to expand them back during read operations.

Sound memory engine 74 has access, shared with the instruction execution unit and the host interface, to blocks of RAM mapped into the address space of memory 72 which implement the TRAM address buffer and TRAM data buffer. These twin buffer memories (TRAM data buffer 84 and TRAM address buffer 85) hold data and address pairs which along with control bits 83, including an align bit 83a and read/write bits 83b, completely specify the activity of the TRAM engine during a sample period. In the way that microprogram memory 80 holds the program executed by the instruction execution unit every sample period, these buffers represent the "program" executed by the TRAM engine every sample

period. Whenever a program compiled for sound effects engine 50 reads or writes to a delay line, it is actually locations in TRAM data buffer 84 that are used as the operands. Ordinarily, address offsets stored in TRAM address buffer 85 are constants initialized by the host processor, but any instruction executed by instruction execution unit 70 can compute a new delay line address by storing its results in the appropriate TRAM address buffer of memory 72.

During each sample period, the TRAM engine runs sequentially through each of the buffers' address/data pairs, so that the buffer contents represent an ordered list of TRAM accesses throughout the entire sample period. During every TRAM memory cycle within the sample period, a TRAM address offset is fetched by the sound memory engine from the TRAM address buffer GPRs, an absolute TRAM memory address is computed from the offset, and a signal value is either fetched from or written to the address paired TRAM data buffer location.

The TRAM data buffers are the source and sink for audio data flowing from TRAM, and each data buffer location is paired one-to-one with an address buffer location; the address buffer holds the address in the TRAM that corresponds to the data. Control bits 83 are a separate array (only writable from the host and not mapped into memory space 72) having a field associated with each data/address pair. Certain of these control bits (control bits 83b -- a two-bit field in this embodiment) specify what type of TRAM operation, e.g., read or write, should take place using the corresponding data and address pair.

By sharing access to the Address and Data buffers with the instruction execution unit, DSP computations and TRAM accesses are decoupled from each other in time and in memory 72 address space. In other words, there is no fixed temporal relationship between when the TRAM engine and instruction execution unit access memory 72. This decoupling is valuable because of the inherent difference in operational speed between the fast instruction execution unit and the relatively slow physical TRAM memory. Decoupling DSP and memory accesses through the shared buffers allows the instruction execution unit to proceed at full speed without regard to the access latencies and pipeline consequences that burden the sound memory engine and its physical memory interface.

Mapping the TRAM buffers into the address space of memory 72 allows the buffer locations to be accessed as operands from instructions for execution unit 70. In this way, data stored in the TRAM can be funneled through the TRAM data buffers and used as input or output operands in DSP programs.

- 5 This transparent memory-mapping means that no special instructions are required for TRAM data access and no extra overhead is incurred waiting for slow TRAM cycles during an access to this data.

Memory 72

- 10 In this embodiment, memory 72 is actually made up of four separate physical memory sections. A more detailed diagram of memory 72 is shown in Fig. 6. As shown in Fig. 6, memory 72 includes 1024 locations, not all of which are used. All 1024 locations in memory 72 are addressable by instruction execution unit 70, however, and thus the entirety of memory 72 can be said to be addressable
15 by a uniform addressing scheme.

- There are four different function areas within memory 72:
input/output space 83 (memory locations 0-40), general purpose memory space 92 (memory locations 100-200), TRAM data and address buffers 84 and 85 (memory locations 200-280: internal TRAM data; locations 280-2A0: external TRAM data;
20 locations 300-380: internal TRAM addresses; and 380-3A0: external TRAM addresses) and hardware registers and constants space 94 (locations 40-60). These functional units do not necessarily correspond with the four physical memories that make up memory 72.

- General purpose memory space 92 includes 256 locations that are
25 32-bits wide for storing general purpose operands for DSP programs run by instruction execution unit 70. This memory space is implemented by a single port memory that operates at four times the instruction rate thereby allowing each instruction to access three input operands and one output operand per cycle. Memory space 94 includes hard-wired frequently used constants along with space
30 for various registers.

As previously stated, TRAM data and address buffers 84 and 85 are ports to and from internal and external TRAMs 76 and 79. Each data buffer

location is paired with an address buffer location. The data buffer holds the TRAM data that was read at, or will be written to, the TRAM address contained in the corresponding address buffer location. Like general purpose memory space 92, these buffers are implemented by a register file memory and appear in the DSP
5 program's operand space. Although they appear in separate parts of memory 72's address space, the data and address portions of the TRAM buffers are accessed in parallel by TRAM engine 74.

In this embodiment, the TRAM buffers are not a full 32-bits wide, but instead are limited to 20-bits. These 20 bits accommodate both the 20-bit
10 TRAM addresses and 20-bit decompressed TRAM data in a single, shared internal RAM. The 20 data bits are left-justified into the most significant position (31) of the 32-bit data path. The 20 address bits, however, are justified into the next significant bit (30) of the data path with the MSB read as a 0. This guarantees that addresses always appear as positive quantities to the instruction execution unit.

Implementation of Delay Lines

Circular addressing refers to the ability to implement delay lines using memory pointers which recirculate through a region of memory. A delay line's read pointer follows its write pointer around the circle and the pointers
20 advance at the sample rate; the distance in memory between the read and write pointer is the delay line length. Circular addressing is also known as modulo addressing since all address calculations are performed modulo the region size, which is typically a power of two for simplicity of implementation.

Even though the read and write pointers are referred to as addresses, in effects processor 50 they are not implemented as physical TRAM memory
25 addresses but as address offsets. The initial value of the offsets can be fixed by an assembler and associated driver software as would be understood by a person of ordinary skill in the art. To get the physical memory address for each read and write operation the offsets are summed with the delay base address counter. This
30 base address counter decrements at the sample rate; thus the sums of the read and write offsets and the base address also decrement at the sample rate, and if the write

offset is smaller than the read offset, the delay length will be the difference of the two as the read pointer chases the write pointer downward through memory.

This is described below and illustrated in Figs. 7A-7C for a four sample delay line implemented in a hypothetical 32-sample TRAM. The numbers highlighted in the circular memory show which values in the delay line are currently "active," waiting to be retrieved. The write offset of the delay line is 0, and the read offset is 4 -- these values remain constant. Initially, the base address (BA) is at 0 (Fig. 7A), so the physical address of the write operation is also 0, where the value 7 is written.

After four sample periods, the base address counter BA has decremented, modulo the TRAM size of 32, from 0 down to 28 (Fig. 7B). This is where the newest value of -2 will be written to the head of the delay line, where the base address value of 28 is summed with the write offset of 0. The fixed read offset of 4 is also added to the base address, modulo 32, to give a physical read address of 0, where it retrieves the value 7 which was written 4 samples previously.

The read pointer leaves values in its wake which are now inactive and can be discarded. As the base address counter begins to decrement back towards 0 it will begin reusing previously written-to memory (Fig. 7C), as the write pointer overtakes discarded values.

This example of delay line implementation in audio signal processor 10 shows that, unlike many other DSPs, circular delay line addressing is not performed modulo the size of the delay line, but modulo the size of the physical address space (which in this miniaturized and simplified example is 32 locations). This means that multiple delay lines of different lengths can occupy the same physical address space and still use the same modulo arithmetic. An example of this is shown in Fig. 7D. In Fig. 7D, the original delay line of Fig. 7A which writes at offset 0 and reads at offset 4 is joined by 3 other delay lines with the following specifications:

2: [Write@5, Read@10],

3: [Write@14, Read@15]

4: [Write@16, Read@28]

For the sake of complete generality, the example uses a Delay Base Address counter value of 2.

As can be seen in Fig. 7D, the four delay lines, each with distinct read and write offsets relative to the Delay Base Address counter, reside simultaneously in TRAM and simply chase each other around the modulo circle formed by the physical address space. Note also that the 4 delay lines do not fill up all of the available memory space: a memory gap exists unused between the operation of a Read@10 and a Write@14, and similarly, there is unused memory past the last Read@28. These areas are available for additional delay lines, or alternatively, the space leftover can be used to allow the delay lines which Read@10 and Read@28 to have their read pointers modulated, that is, expanded and contracted dynamically, without interfering with the other delay lines.

Audio effects processor 50 also makes it easy to modulate delay line addresses, since addresses are stored in memory 72 in the TRAM address buffer 84 and can be accessed just like any other arithmetic operand. The length of a delay line is modulated by computing new addresses for the read and/or the write pointer of the delay line, although it is most typical to change only the read pointer.

The length of a delay line is the distance in samples from the write address to the read address, so to properly modulate a delay line the new computed delay length must be summed with the delay line's write address before storing it in the TRAM address buffer. Since the write address stored in the buffer is eventually used again as an offset from the delay base address counter to get the physical TRAM address, the modulated delay length really represents an additional offset from an offset from the delay base address counter. Since the delay line's write address is already available in the TRAM address buffer, the read offset can be computed with an instruction of the form:

$$\text{Read Address} = \text{Write Address} + (\text{Mod Function} \times \text{Max Delay Length})$$

where Max Delay Length is the declared maximum length of the delay line, and the Mod Function is a control signal that varies from 0 to 1.0.

Current VLSI processes (0.35 micron) do not cost-effectively support the amount of internal memory that is needed to implement the entire delay line capacity required for some complex sound effect applications like reverb and 3D audio (from .5Mbit to 2Mbit). On the other hand, prudent use of PCI bus 49 does not support the bandwidth to external host memory required by some applications (from 64 to 128 accesses/sample period). Audio effects processor 50 compromises by offering access to both these types of memory: high-bandwidth access to 8K samples of internal TRAM 76 (ITRAM) implemented in on-chip, single-port static memory, and low-bandwidth access to up to 1M or larger samples of external TRAM 78 (XTRAM) implemented across PCI bus 49 in host DRAM memory.

These two memory spaces, ITRAM and XTRAM are logically and physically distinct and their address spaces overlap. For example, address 0x0000 in ITRAM is not the same location as address 0x0000 in XTRAM. Thus, delay lines must be implemented either in one memory space or the other.

Avoidance of Read/Write Hazards

There is no double-buffering of TRAM data buffers 84. Thus, situations can arise where the microprogram accesses data in the buffers which does not strictly belong to the current sample period. In the case of read operations, a microinstruction can read a datum from the buffer that was actually fetched from TRAM during the previous sample period. In the case of write operations, a microinstruction may write a datum to the TRAM buffer that will not actually be flushed out to TRAM until the next sample period.

These situations arise from the fact that sound memory engine 74 is not tightly coupled to instruction execution unit 70. There is an arbitrary temporal relationship between the accesses performed by instruction execution unit 70 and sound memory engine 74 on the same datum in the buffer. The relationship is established by the relative positioning of the microinstruction word and the TRAM buffer location and the relative access rates of the execution unit and sound memory engine. The relationship is determined by a microprogram loader (e.g., software driver) as would be understood by a person of skill in the art.

An example of a recall hazard from a TRAM read operation is shown in Fig. 8A. Fig. 8A shows two TRAM read operations, one from address offset 100 and the other from offset 110. The Base Address Counter, which decrements by 1 each sample period, is added to the TRAM offsets stored in the TRAM address buffer during every access to form a physical TRAM memory address. For the sake of simplicity, this example assumes that the TRAM and instruction execution engines are sequencing at the same rate through the buffer and microprogram memories, respectively.

The fetch operation occurs as the TRAM engine encounters the TRAM read offset 100. It initiates a memory fetch from physical address 101 and after a TRAM access delay time the data from 101 appears in the data buffer. It is subsequently read from the buffer by the execution of a microinstruction. In this case, the fetch preceded the read, and the microinstruction received its data from TRAM address 101, offset from the base address of 1 by 100, just as it expected.

Some number of instruction cycles later, a microinstruction seeking data from offset 110 reads its data from the buffer location nominally filled with data from 110. In this case, however, the read operation precedes the TRAM fetch, which won't occur for another few cycles, since the TRAM offset 110 is stored in the address buffer a few locations past the microinstruction word. Instead of receiving data from TRAM address 111 (the offset 110 plus the current Base Address Count of 1), the microinstruction receives data that was fetched during the previous sample period. At that time the Base Address Counter was 2, so the datum in the buffer that is read is actually from TRAM address 112, delayed exactly one sample period too long.

This off-by-one error occurs whenever microinstruction reads precede TRAM buffer fetches. A 1-bit address offset (align bit 83a) along with align control circuit 87 provides a solution to this problem by canceling out the off-by-one error whenever microinstruction reads precede TRAM engine fetches. Using the same example as before, this solution is illustrated in Fig. 8B. As shown in Fig. 8B, the TRAM buffer structure has been augmented to include a field containing address align bit 83a, which is paired 1-to-1 with the TRAM addresses. Align control

circuit 87 examines align bit 83a and provides an input to address generator 88 that is used in determining the physical TRAM address.

A problem symmetrical to read hazards can exist for TRAM write operations. This situation is illustrated in Fig. 8C. When microinstructions write results to TRAM, the result data is buffered just as for read operations, waiting for the TRAM engine to flush the result out to memory. There is no problem if microinstructions that write to the TRAM are executed before the flush operation takes place. The difficulty arises when the TRAM offset and data are stored in buffer locations preceding the microinstruction word. In these cases the TRAM engine is actually flushing data that was written by the microcode during the previous sample period; the effective result is that the delay line includes one sample period of delay beyond its nominal length.

To address this issue, sound memory engine 74 ensures that during the subsequent sample period that the data is flushed out to exactly where it should have ended up. If the base address is counting down, this means "1" is added to the offset to compute the physical address to which we should have flushed the data during the previous sample period. Again, and as shown in Fig. 8D, align bit 83a can be used by align control circuit 87 to cancel this off-by-one error for write operations.

The difference in this case is that the align bit signifies a "+1" offset going into the TRAM address engine instead of the "-1" used for read operations. This is easily accomplished by taking the read/write bits associated with each address (not shown) into account.

One strategy for setting the align bits is to use a microcode loader (software driver), which keeps track of the location of all offsets in the TRAM address buffer as well as all instructions which use the TRAM data corresponding to those offsets. This "loader" may be running either at microcode download time or might really just be a "locator" which only needs to run at microcode build time. With information and knowledge about the relative execution rate of the microcode and TRAM engine sequencers, the loader can determine which TRAM data will be read from the buffer before it is actually fetched from memory or written after it has been flushed. It will flag such situations by setting the align bit corresponding

to the TRAM offset associated with that data. The align bit is then fed into TRAM address engine 88 and used as an extra offset of ± 1 when forming the physical TRAM address.

Another approach is to use hardware in sound memory engine 74 to automatically generate the align bit offsets. The microprogram sequencer restarts on the same cycle during which the TRAM engine sequencer is restarted at the beginning of the buffer and the base address count is incremented. At this time, the align bit is cleared for every read offset, and set for every write offset. As the microcode accesses the TRAM data buffer for read or write operations, the align bit for the corresponding offset is set for every read offset and cleared for every write offset. Thus, "misaligned" TRAM fetch or flush operations will automatically have used the correct sense of the align bit. This approach is less preferred because other measures must be taken to ensure that it is not possible to slip a microinstruction access to the TRAM data buffer into the pipeline delay between address usage and data fetch. Without such assurance, the microinstruction access may not have the opportunity to update the alignment bit even though it was accessing last sample period's buffer data, and the physical address may never be correct.

20 Initialization of TRAM Data and Address Buffers 84 and 85

When delay lines are used, TRAM read and write operations start as soon as power is received by effects processor 50. Write operations do not pose a problem, but unless TRAMs 76 and 78 are properly initialized, read operations will initially produce noise. Thus, it is important to initialize the TRAMs. In one embodiment, these memory spaces are initialized by writing 0s to each location before executing a program in sound processor 50. This can be a time consuming task, however. Thus, in other preferred embodiments, dedicated hardware circuits are employed to perform such an initialization routine. As would be understood by a person of skill in the art, a variety of such circuits can be designed and employed in the present invention.

Other Embodiments

The above is a full, detailed description of one specific embodiment of the present invention. It should be understood that many specifics of the invention described above can be changed without departing from the concept of the invention. For example, word size, address field length, memory size, number of control bits 83 used and other implementation-specific items can vary in other embodiments. Also, while effects processor 50 was described as an audio effects processor that performs operations on audio signals in a sound processing system, a person of ordinary skill in the art will realize that effects processor 50 can also process radar signals, seismic signals or any other time series data in any signal processing system.

Many other equivalent or alternative methods of implementing audio effects processor 50 will be apparent to those skilled in the art. For example, it is possible for sound memory engine 74 to implement table lookup operations using non-circular indexed addressing in internal and external TRAM memories 76 and 78. Tables of data can be loaded into TRAM by host processor 18 and then accessed like static arrays whose indices are computed by instruction execution unit 70. The sound memory engine can distinguish delay line accesses from table-lookup accesses by the setting of additional TRAM control bits associated with each address. Table-based addressing can be useful for implementing non-linear functions, wavetables or excitation functions for synthesis, block-based coefficient storage, or for many other applications. Note that although "table lookup" implies a read operation, table-based addressing applies equally to write operations, so it is possible to have audio effects processor 50 compute table values in real time instead of performing such calculations with the host.

Table-based addressing is distinguished from delay line addressing by the use of a static table base address register in place of the constantly decrementing delay base address counter. In the same way the delay line accesses are always made with offsets relative to the delay base address counter, table lookups are always made relative to a table base address register. The table base address register is initialized by the host and mapped into the address space of memory 72.

If necessary, the table base address can be read (but not written) from DSP programs like any other instruction operand.

Initialization of the table base address register is typically performed by host 18 at system startup time. The address stored there partitions physical

5 TRAM into a circular region and a table-lookup region. All delay line addresses are computed modulo the table base address instead of modulo the physical memory size, so that delay line operations automatically take place within the memory region from 0-[Table Base]. In this way the table base address register also serves in effect as a "delay limit" register.

10 In the same way that circular delay addressing treats the TRAM address buffer elements as offsets from the delay base address counter, table-based addressing uses these addresses as offsets relative to the table base address register. To compute the physical TRAM address, the TRAM address buffer contents are added to the table base address register. By calculating a table index with the
15 proper scaling and storing it as an offset in the TRAM address buffer, DSP programs can perform computed table lookups.

All programs loaded into audio effects processor 50 use the same table base address register, whose value is typically established by the host at system initialization time. If more than one program requires table lookup storage,
20 the individual tables belonging to different programs are partitioned from within the global table lookup region. DSP programs can declare their own private table offset registers using ordinary GPRs whose contents represent an offset from the table-based address register, and whose location and content will be properly relocated by the program loader. Then these private table offset register GPRs are
25 summed with the computed table index before performing the table look-up operation in TRAM memory.

Typically table-lookup offsets will create addresses which are inside the table lookup region of the TRAM, but since the sound memory engine performs no address limiting or trapping it is possible for specialized applications to generate
30 addresses outside this region. As for delay line addressing, the operand data path alignment of the TRAM address buffer guarantees that all TRAM table offsets are positive quantities, so to generate addresses below the table lookup region the

program must generate addresses which wrap around the entire TRAM address space (2^{20} locations in one embodiment) and back into the delay line region.

In addition to implementing table look-up operations, other variations of the present invention are also possible. For example, while sound processing engine 40, sound processor 50 and interface unit 60 were all described as being part of a single chip sound processor, the functionality of these elements can be performed from separate ICs in other embodiments. Similarly, it is possible to develop an architecture in which instruction execution unit 70 and sound memory engine 74 are implemented on separate chips. Also, sound processor 30 can be included directly on a computer motherboard rather than a dedicated multimedia card 32, and in another data is not compressed and expanded when it is transferred between memory 72 and TRAM memories 76 and 78. These equivalents and other alternatives are intended to be included within the scope of the present invention.

**WHAT IS CLAIMED IS:**

- 1 1. A signal processor comprising:
 - 2 (a) a first memory including a first address space;
 - 3 (b) an instruction execution unit that processes signal
 - 4 processing instructions including instructions to read from and write to said first
 - 5 address space; and
 - 6 (c) a signal memory engine that operates independent of
 - 7 said instruction execution unit to read from and write to said first address space and
 - 8 perform address arithmetic and memory accesses to a second memory as necessary
 - 9 to implement signal data storage.
- 1 2. The signal processor of claim 1 wherein said instruction
- 2 execution unit executes instructions at a first execution rate and wherein said signal
- 3 memory engine executes operations at a second execution rate different than said
- 4 first execution rate.
- 1 3. The signal processor of claim 1 wherein said first address
- 2 space includes a first memory location and wherein said signal memory engine and
- 3 instruction execution unit can independently access said first memory location in a
- 4 non-fixed temporal relationship.
- 1 4. The signal processor of claim 1 wherein said signal memory
- 2 engine performs arithmetic and executes memory accesses to said second memory
- 3 based on values stored in said first address space.
- 1 5. The signal processor of claim 1 wherein said signal data
- 2 storage comprises audio data storage.
- 1 6. The signal processor of claim 5 further comprising a sound
- 2 processing engine that synthesizes sounds.



1 7. The signal processor of claim 6 wherein output from said
2 sound processing engine representing synthesized sounds is stored in a second
3 address space of said first memory and wherein instructions processed by said
4 instruction execution unit read from and write to said second address space.

1 8. The sound processor of claim 1 further comprising an
2 instruction memory coupled to said instruction execution unit, said instruction
3 memory comprising a microprogram containing a specified set of signal processing
4 instructions to be executed by said instruction execution unit.

1 9. The signal processor of claim 8 wherein said instruction
2 execution unit executes a single pass through its entire microprogram each sample
3 period in an order in which the instructions appear in said microprogram without
4 branches, loops or subroutine calls.

1 10. The signal processor of claim 9 wherein said instruction
2 execution unit provides for conditional execution of instructions by ignoring
3 execution results of selected instructions based upon values of a condition code for
4 said selected instructions.

1 11. The signal processor of claim 8 wherein said instruction
2 memory is a different memory than said first memory.

1 12. The signal process of claim 1 wherein said first memory is a
2 randomly addressable memory.

1 13. The signal processor of claim 12 wherein said randomly
2 addressable memory comprises a plurality of separate, physical RAM memories.

1 14. The signal processor of claim 12 wherein said first memory
2 further includes a second address space for storing values representative of sound
3 input and sound output.



1 15. The signal processor of claim 12 wherein said first memory
2 further includes a third address space for storing frequently used coefficients,
3 variables and intermediate values.

1 16. The signal processor of claim 15 wherein said instruction
2 execution unit uses a uniform addressing scheme to access said first, second and
3 third address spaces of said first memory.

1 17. The signal processor of claim 5 wherein said audio data
2 storage is accessed using circular addressing to implement audio delay lines.

1 18. The sound processor of claim 5 wherein said audio data
2 storage is accessed using table-based addressing.

1 19. A single chip sound processor comprising:
2 (a) a first memory including a first address space and a
3 second address space;
4 (b) a second memory that stores sound processing
5 instructions;
6 (c) a sound processing engine that synthesizes sounds and
7 outputs data representative of said synthesized sounds to said second address space;
8 (d) an instruction execution unit configured to execute said
9 sound processing instructions stored in said second memory, including instructions
10 to reads from and write to said first and second address spaces; and
11 (e) a sound memory engine that operates independent of
12 said instruction execution unit to read from and write to said first address space and
13 perform address arithmetic and memory accesses to a second memory as necessary
14 to implement audio data storage.

1 20. The sound processor of claim 19 wherein said instruction
2 execution unit executes instructions at a first execution rate and wherein said sound



memory engine executes operations at a second execution rate different than said first execution rate.

21. The sound processor of claim 19 wherein said first address space includes a first memory location and wherein said sound memory engine and instruction execution unit independently access said first memory location in a non-fixed temporal relationship.

22. A computer processing system comprising:

- (a) a bus;
- (b) a host processor communicatively-coupled to said bus;
- (c) a second memory; and
- (d) a sound processor communicatively-coupled to said bus, said sound processor comprising:
 - (i) a first memory including a first address space;
 - (ii) an instruction execution unit that processes sound processing instructions including instructions to read from and write to said first address space; and
 - (iii) a sound memory engine that operates independent of said instruction execution unit to read from and write to said first address space and perform address arithmetic and memory accesses to a second memory as necessary to implement audio data storage; and
 - (iv) a host interface unit that interfaces said second processor to said bus.

23. The computer processing system of claim 22 wherein said host processor can read from and write to said first address space.

24. The method of implementing an audio delay line in a sound processor, said method comprising the steps of:



- 3 (a) executing a sound processing instruction in an
4 instruction execution unit of said sound processor to store data representative of a
5 sound to be delayed in a first memory;
6 (b) retrieving, independent of step (a), said data from said
7 first memory with a sound memory engine of said sound processor and writing said
8 data to a second memory; and
9 (c) performing address arithmetic and subsequent memory
10 accesses to said first and second memories with said sound processor to implement
11 the audio delay line.

1 25. The method of claim 24 further comprising the step of
2 executing an instruction in said sound processor to write data stored in said first
3 memory by said sound memory engine to a different location in said first memory
4 for output from said sound processor.

1 26. A method of operating a computer processor, said method
2 comprising the steps of:
3 (a) synthesizing an audio signal with a sound synthesis
4 engine of said computer processor;
5 (b) outputting said audio signal to a sound effects engine of
6 said computer processor;
7 (c) storing said outputted audio signal in a first location of
8 said memory coupled to said sound effects engine;
9 (d) executing a first sound processing instruction in an
10 instruction execution unit of said sound effects engine to transfer said stored audio
11 signal to a second location of said memory, said second location being paired with
12 an address to a location in a sound delay memory; and
13 (e) performing, independent of step (d), address arithmetic
14 and memory accesses to said sound delay memory with a sound memory engine
15 independent of step (d) to implement an audio data storage.

1 27. The method of claim 26 wherein step (e) comprises the step
2 of, after storing said audio signal in said sound delay memory for at least one
3 sample period, writing said audio signal to a third memory location of said first
4 memory.

1 28. The method of claim 27 further comprising the step of
2 executing a second sound processing instruction in said instruction execution unit to
3 write data stored in said third location of said first memory by said sound
4 processing engine to a fourth location of said first memory for output from said
5 sound effects engine.

