



(51) International Patent Classification:
G06F 15/78 (2006.01)

(21) International Application Number:
PCT/US2018/052346

(22) International Filing Date:
24 September 2018 (24.09.2018)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/566,146 29 September 2017 (29.09.2017) US

(71) Applicant: **KNOWLES ELECTRONICS, LLC**
[US/US]; 1151 Maplewood Drive, Itasca, Illinois 60143
(US).

(72) Inventors: **RUB, Leonardo**; c/o Knowles Electronics,
LLC, 331 Fairchild Drive, Mountain View, California
94043 (US). **CLARK, Brian**; c/o Knowles Electronics,
LLC, 331 Fairchild Drive, Mountain View, California 94043
(US).

(74) Agent: **DANIELSON, Mark J.** et al.; Foley & Lardner
LLP, 3000 K Street NW, Suite 600, Washington, District of
Columbia 20007 (US).

(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ,
CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP,
KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,
SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ,
UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,
EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,

(54) Title: MULTI-CORE AUDIO PROCESSOR WITH LOW-LATENCY SAMPLE PROCESSING CORE




FIG. 2

(57) Abstract: A multi-core audio processor includes a data protocol interface configured to receive a stream of audio data, a plurality of data processing cores including a single sample processing core and a block data processing core, an audio fabric block configured to route samples of the stream between the data protocol interface and the plurality of data processing cores. The single sample processing core includes an execution unit configured to execute one or more low latency instructions for performing computations for the samples.

TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

MULTI-CORE AUDIO PROCESSOR WITH LOW-LATENCY SAMPLE PROCESSING CORE

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of and priority to U.S. Provisional Patent Application No. 62/566,146, filed September 29, 2017, the entire contents of which are incorporated herein by reference.

BACKGROUND

[0002] It is desirable to minimize the latency of processing tasks in digital signal processing. For example, when processing a stream of audio data, it is desirable to process samples of the stream as quickly as possible before overrun or underflow occurs. Certain applications, such as asynchronous sample rate conversion, phase alignment of multiple streams, ambient noise cancellation, and speaker excursion protection are particularly sensitive from a latency perspective.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] FIG. 1 is a block diagram of an environment of a multi-core audio processor.

[0004] FIG. 2 is a block diagram of a single sample processing core of the multi-core audio processor depicted in FIG. 1.

[0005] FIG. 3 is a block diagram of various components of an execution unit of the single sample processing core shown in FIG. 2.

[0006] FIG. 4 is a chart depicting the contents of a vector register file of the execution unit shown in FIG. 3.

[0007] FIG. 5 is a schematic diagram of operations performed via the execution unit shown in FIG. 3.

[0008] FIG. 6 is an operational diagram of an asynchronous sampling rate conversion filter.

[0009] FIG. 7 is a block diagram depicting the execution of instructions by an execution unit to compute a sine and cosine of an angle and an exponential function of a magnitude.

[0010] FIG. 8 is a block diagram depicting the execution of a single instruction by a low-latency execution unit to compute a sine and cosine of an angle and an exponential function of a magnitude.

[0011] FIG. 9 is a block diagram depicting various operations of a low latency execution unit computing a quotient.

DETAILED DESCRIPTION

[0012] The present embodiments will now be described in detail with reference to the drawings, which are provided as illustrative examples of the embodiments so as to enable those skilled in the art to practice the embodiments and alternatives apparent to those skilled in the art. Notably, the figures and examples below are not meant to limit the scope of the present embodiments to a single embodiment, but other embodiments are possible by way of interchange of some or all of the described or illustrated elements. Moreover, where certain elements of the present embodiments can be partially or fully implemented using known components, only those portions of such known components that are necessary for an understanding of the present embodiments will be described, and detailed descriptions of other portions of such known components will be omitted so as not to obscure the present embodiments. Embodiments described as being implemented in software should not be limited thereto, but can include embodiments implemented in hardware, or combinations of software and hardware, and vice-versa, as will be apparent to those skilled in the art, unless otherwise specified herein. In the present specification, an embodiment showing a singular component should not be considered limiting; rather, the present disclosure is intended to encompass other embodiments including a plurality of the same component, and vice-versa, unless explicitly stated otherwise herein. Moreover, applicants do not intend for any term in the specification or claims to be ascribed an uncommon or special meaning unless explicitly

set forth as such. Further, the present embodiments encompass present and future known equivalents to the known components referred to herein by way of illustration.

[0013] Referring generally to the figures, described herein is a multi-core audio processor including a single sample processing core. The single sample processing core is tailored for performing specific tasks for sample-based digital signal (e.g., audio) processing with minimal latencies. In this regard, the single sample processing core includes execution units configured via various sets of hardware instructions. These execution units implement datapath elements (e.g., operation units) defined by the hardware instructions via a finite state machine of the single sample processing core. The datapath elements are configured to minimize the latency in the implementation of various signal processing applications (e.g., asynchronous sample rate conversion, phase alignment of multiple streams, ambient noise cancellation, speaker excursion protection). In a first example, a first execution unit is configured via a first set of hardware instructions to include sine and cosine computation units communicably coupled to a first location of a vector register and an exponential function computation unit communicably coupled to a second location of the vector register. The first execution unit computes sine and cosine values of an angle stored at the first location and an exponential value of a magnitude stored at the second location during execution of a single instruction to minimize the number of execution cycles required for the computation of the sine, cosine, and exponential values.

[0014] In another example, a second execution unit is configured via a second set of hardware instructions to compute a quotient of an elapsed time and an output sample time increment. The elapsed time may be in units of input samples. For example, the quotient may be used to identify a number of output samples to generate per input sample of audio data received via an I/O interface of the single-sample processing core. A fixed point division computation unit may be employed to compute a quotient of a fractional part of the elapsed time and a portion of the output sample time increment with minimum latency. The quotient may be used to identify a number of output samples to compute per input sample received to provide zero overhead loops when executing an output stage of a filter. In some embodiments, the first and second execution units are employed during an implementation of

an asynchronous sampling rate conversion (ASRC) algorithm to minimize the latency thereof.

[0015] In another example, a third execution unit may be configured by a third set of hardware instructions to include a plurality of operation (e.g., multiplication) units, a vector register file, and a set of state registers. In some embodiments, there are a plurality of hard-wired routes coupling various locations of the vector register file to the operation units. The set of state registers store values assigning specific locations within the vector register file to specific ones of the operation units. The third execution unit is configurable to execute a set of program instructions. The program instructions may include configurable element selection depending on a value of an immediate field of an instruction. In an example, if the most significant bit of the immediate field is zero, the remaining bits of the immediate field specify a set of hard-wired routes to be enabled for directing data stored in the vector register file to the operation units. If the most significant bit of the immediate field is one, the remaining bits of the immediate field specify one of the state registers, and an assignment value stored therein is used to assign locations of the vector register file to specific operation units. The use of the state registers eliminates the need to re-load routing registers during multiple executions of program instructions, thereby minimizing latency incurred in performing the operations via the operation units. Additionally, the inclusion of the hard-wired routes diminishes the number of state registers necessary to reduce hardware costs.

[0016] Referring now to FIG. 1, an environment 100 including a multi-core audio processor 102 is shown, according to an example embodiment. Generally, the multi-core audio processor 102 is configured to receive input streams of data (e.g., audio data) from a plurality of input devices, perform various processing tasks (e.g., noise cancellation, filtering, etc.) on the input streams via a plurality of processing cores 110, 112, and 114, and transmit output streams of data to at least one output device 118.

[0017] The input devices 116 may communicate input streams to the multi-core audio processor 102 having a variety of different characteristics. For example, a first input device 116 may input a physical stream of a first data type (e.g., pulse code modulation-encoded data) and a second input device 116 may input a physical stream of a second data type (e.g.,

pulse density modulation-encoded data). To facilitate the processing of such a variety of data types, the multi-core audio processor 102 includes an audio transport block 104. The audio transport block 104 includes a plurality of different types of interfaces (e.g., audio interfaces) configured to perform operations (e.g., serialization and deserialization, filtering, etc.) on streams after their reception from the input devices 116 and prior to their transmittal to the output devices 118 (e.g., via data ports). For example, in one embodiment, the audio transport block 104 includes a plurality of PCM interfaces and a plurality of PDM interfaces configured to serialize and de-serialize constituent channels of the streams received via the input devices 116 and processed via the processing cores 110, 112, and 114.

[0018] In various embodiments, streams input to the multi-core audio processor 102 each include a plurality of coherent channels having the same sampling rate. To facilitate low latency processing of the streams, it is desirable to maintain this coherency when storing samples of the streams in queues for processing. In this regard, the multi-core audio processor 102 further includes a plurality of transport buffers 106. In some embodiments, the multi-core audio processor 102 includes a transport buffer for each audio interface included in the audio transport block 104. In some embodiments, each transport buffer includes a plurality of asynchronous first-in-first-out (FIFO) buffers. Groupings of the FIFO buffers may be coupled to one another so as to synchronously store data of constituent channels of various streams to maintain the coherency of the constituent streams. Each stream may have an associated identifier. The identifier for a stream may be correlated with the audio interface used to communicate data of the stream. For example, in one embodiment, the audio transport block 104 include six PCM interfaces. If a stream is received by the sixth PCM interface, the stream identifier may identify that interface (e.g., PCM 6). In various embodiments, the multi-core audio processor 102 also includes an audio fabric block 108 configured to store such identifiers in accordance with a logical address map. For example, each identifier may be stored in a location that identifies the grouping of FIFO buffers storing data of a corresponding stream. This way, upon receipt of an identifier (e.g., from one of the processing cores 110, 112, and 114), the audio fabric block 108 can route data (e.g., via routing hardware including a set of multiplexers enabled via the logical address map) to a requesting processing core. In some embodiments, the audio fabric block 108 also includes sample counters configured to count numbers of samples received at the audio transport

block. Additionally, the audio fabric block 108 may also include a presentation timer configured to cause the transport buffers 106 to provide presentation times to the processing cores 110, 112, and 114 when the processing cores read data via the audio fabric block.

[0019] The processing cores 110, 112, and 114 may be digital signal processing cores having different processing capabilities. In the example shown, the processing core 110 is a single sample processing core and the processing cores 112 and 114 are frame processing cores. As such, the processing core 110 includes various hardware components designed specifically to perform sample-based (e.g., the application of filter banks, sample rate conversion, etc.) processing tasks on input samples received via the audio transport block 104. The details of such hardware components are described in greater detail herein with respect to FIGS. 2-9. The processing cores 112 and 114, on the other hand, include various hardware components designed to perform block data processing tasks on sets of samples (e.g., frames) of the streams received via the audio transport block 104. Thus, by the inclusion of multiple processing cores optimized to perform different types of data processing, the multi-core audio processor is able to perform multiple tasks simultaneously to maximize the speed with which various functions are carried out.

[0020] In some embodiments, the multi-core audio processor 102 further includes non-digital signal processing cores, such as a main controller processor configured to arbitrate undedicated resources (e.g., memory) between the processing cores. The main control processor may be of an ARM architecture. Additionally, a system controller processing core may be configured to handle other tasks (e.g., system wakeup, power down, etc.). In some embodiments, the processing core 110 is directly connected to the audio fabric block 108 via a native bus of the audio fabric block 108. The other processing cores may be communicably coupled to the audio fabric block 108 by a processor interconnection structure. For example, the processor interconnection structure may couple each of the processing cores of the multi-core audio processor 102 to one another and be coupled to the audio fabric block 108 via a communications bridge.

[0021] In various embodiments, the multi-core audio processor 102 includes a dedicated memory (e.g., random access memory) for each processing core contained there. For

example, the dedicated memory may store program instructions and data serving as inputs to the program instructions executed by the processing cores 110, 112, and 114. Additionally, the multi-core audio processor 102 may also include a pool of undedicated memory. The main controller processor may update allocation logic of a central memory switch to change how the undedicated memory is allocated amongst the processing cores.

[0022] Referring now to FIG. 2, a schematic diagram of the single sample processing core 112 of FIG. 1 is shown. As shown, the single sample processing core 112 includes an instruction memory interface 202 (e.g., associated with an address bus and an instruction bus) in communication with an instruction memory (e.g., a portion of a dedicated memory of the single sample processing core 112) and a data memory interface 206 in communication with a data memory (e.g., another portion of the dedicated memory). The single sample processing core 112 further includes an instruction fetch and decode unit 208 configured to decode program instructions received via the instruction memory interface 202.

[0023] In various embodiments, the single sample processing core 112 includes a set of base execution units 216 (e.g., load and store units, arithmetic logic units, etc.) of various types configured to execute base program instructions for base data processing tasks (e.g., fixed and floating point scalar and vector arithmetic operations, etc.). Such base program instructions may have base execution pipelines 212 having an associated pipeline control breaking up instruction lifecycles into parallel segments to maximize data throughput in the performance of the base data processing tasks. The base execution units 216 execute the base program instructions to perform tasks (e.g., scalar and vector operations, etc.) on data loaded from memory via the data memory interface 206.

[0024] In addition to the base execution units 216, the single sample processing core 112 further includes a plurality of low latency execution units 214. For example, in various embodiments, datapath elements of the low latency execution units 214 are generated via a firmware-controlled state machine of the single sample processing core 112. Firmware controlling the datapath elements includes a set of hardware instructions chosen to provide the low latency execution units 214. The datapath elements include various operation units and task-specific registers (e.g., vector register files, state registers, etc.) enabling the

execution of a set of low latency audio processing instructions. The low latency audio processing instructions may have low latency execution pipelines 210 having an associated pipeline control breaking up lifecycles of the low latency instructions into parallel segments to maximize data throughput. As described herein, the low latency instructions executed via the low latency execution units 214 are designed to minimize the latency with which the single sample processing core 112 performs various tasks in processing of samples of the streams received via the audio transport block 104. In other words, via the incorporation of the low latency execution units 214, the single sample processing core 112 is able to execute program instructions specifically designed to lower latency in crucial data processing tasks.

[0025] The single sample processing core 112 further includes a processor interconnect interface 204. The processor interconnect interface 204 is communicably coupled to a system bus of the multi-core audio processor 102. In some embodiments, via the processor interconnect interface 204, the single sample processing core 112 is in communication with the audio fabric block 108. As such, input samples of various streams received via the audio transport block 104 may be communicated between the transport buffers 106 and the single sample processing core 112 via the SSP processor interconnect interface 204. In some embodiments, the single sample processing core 112 also includes a queue interface directly coupled to the audio fabric block 108 via a native bus thereof. Samples of audio data received from the audio fabric block 108 may be stored in system memory associated with the single sample processing core 112 via load and store units of the single sample processing core 112. Once stored in memory in association with an address, the sample may be retrieved via execution of program instructions (e.g., base instructions, low latency instructions, etc.) to perform various operations thereon.

[0026] Referring now to FIG. 3, a schematic diagram of various datapath elements associated with a low latency execution unit 214 is shown. As described herein, the single sample processing core 112 is configured to include such data path elements via hardware instructions decoded by a firmware controlled state machine. The datapath elements described with respect to FIG. 3 are exemplary only. Various elements are left out for the purposes of illustration.

[0027] As shown, the datapath elements include a low latency execution unit 214 including operation units 300 arranged to facilitate execution of low latency program instructions. The operation units 300 may include hardware-implemented computation units in communication with a register file to facilitate the retrieval of data from the register file, the performance of computations on the retrieved data, and the re-storage of the computed data into the register file based on inputs of low-latency instructions executed via the low latency execution unit 214. In the example shown, the register file is embodied as a vector register file 304 including a plurality of series of individual multi-bit (e.g., thirty two bit) registers. Data stored within address ranges of a system memory may be loaded into the vector register file (e.g., via the data memory interface 206) in unison to reduce latency of computations performed on the data.

[0028] Additionally, the datapath elements also include a set of state registers 302 storing information used to control the operation units 300 in the execution of the low latency instructions. In various embodiments, the state registers set control bits associated with the operation units 300 based on a value an immediate field input of a low latency instruction being executed via the low latency execution unit 214. As described herein, the state registers 302 facilitate flexibility in terms of operations performed on the data stored in the vector register file 304.

[0029] FIG. 4 is a chart 400 depicting the contents of the vector register file 304. As shown, the vector register file 304 includes eight vector registers V0:V7. Each vector register includes eight multi-bit registers. In the example shown, the vector register file 304 includes thirty two bit individual registers arranged as shown, such that each vector register V0:V7 may store eight thirty two bit values VX[0:7].

[0030] Referring now to FIG. 5 is a schematic diagram illustrating operations via an execution unit (e.g., a vector arithmetic unit) upon execution of an instruction is shown, according to an example embodiment. As shown, a vector register file 500 (e.g., similar to the vector register file 304 described with respect to FIGS. 3-4) includes a plurality of eight-entry vector registers. The execution unit is configured to execute a vector multiply instruction $Z = \text{vmul}(X, Y)$ via a plurality of multiplication units 502, where X represents a

first vector register of the vector register file 500 (e.g., $X[0:7]$) and Y represents a second vector register of the vector register file (e.g., $Y[0:7]$). As shown, the multiplication units 502 include inputs A and B that are wired to the vector registers identified via the inputs to the vector multiply instruction.

[0031] As such, the $vmul$ instruction performs the following operations $Z[0] = X[0]*Y[0]$, $Z[1] = X[1]*Y[1]$, $Z[2] = X[2]*Y[2]$, etc. Thus, while this instruction achieves acceleration over a scalar execution unit via performing multiple operations simultaneously, the instruction is limited since each element of the first input vector X can only be multiplied by a corresponding element of the second input vector Y .

[0032] A way to further increase flexibility is to add an immediate field to the $vmul$ instruction. The immediate field adds flexibility in assigning vector elements $X[0:7]$ and $Y[0:7]$ to different inputs A and B of the multiplication units 502. For example, one may define an additional vector U including a concatenation of X and Y and assign indices associated with the inputs A and B of the different multiplication units 502. In this example, since there are eight multiplication units 502, there are sixteen indices that must be assigned to the additional vector U . As such, sixty four bits are needed to represent each of the 16 indices, which means that 64 bits are necessary to provide complete flexibility in terms of assignments to the inputs.

[0033] Since the addition of a sixty four bit immediate field to the $vmul$ instruction is not practical, one way to achieve this flexibility is to add a register file R to the datapath of the execution unit. The register file R may include 64-bit registers storing assignment values to the multiplier inputs A and B . Thus, a new multiply instruction $Z = vmulR(X, Y, P)$, may be defined based on the new datapath. In this example, P specifies one of the registers of the register file R and thus identifies a set of assignments for the inputs of the multiplication units 502. However, there may be sets of assignments between the input vectors X and Y and the multiplication units 502 that are particularly common (e.g., such as the assignments shown in FIG. 5). Storing such assignments in the register file R would be wasteful. To prevent wasteful inclusion of additional registers in the register file R , one may define a second instruction $Z = vmulH(X, Y, H)$, where H is an immediate field specifying one of the most

common routes. Thus, when one desires to use one of the common sets of assignments, one executes the vmulH instruction, and when one desires to use an uncommon set of assignments, one executes the vmulR instruction. This pair of instructions minimizes the necessary size of the register file R while adding flexibility over the vmul instruction.

[0034] When programming the multi-core audio processor in a high level language (e.g., C), function calls are used to reduce code size and achieve modularity. While this speeds up software development, the use of such high level language may add latency because variables need to be defined and loaded from memory. For example, referring to the vmulH and vmulR instructions above, the need to load the register file R adds the latency of the multiplication operations. This is illustrated by the following pseudocode:

```
void func1()
{
    // Register variable types:
    // can reside in registers or memory (as automatic variables)
    R P; // Not known which register of R is P, so it must be loaded here
    R Q; // Not known which register of R is Q, so it must be loaded here
    V X;
    V Y;
    V Z;

    .
    .
    .
    Get routing from memory and load P; // Adds to latency
    Z = vmulR(X, Y, P);
    .
    .
    .
    Get routing from memory and load Q; // Adds to latency
    Z = vmulR(X, Y, Q);
    .
    .
    .
}
```

[0035] In other words, the need to load the assignment values from memory into the routing register R adds to the latency of the multiplication operations. This can be avoided by adding a set of state registers S0, S1, S2, etc. to the datapath of the execution unit. For example, in one embodiment, eight sixty four bit state registers are added to replace the register file R. Additionally, the instructions vmulR and vmulH are replaced with a new instruction having the syntax $Z = \text{vmulS}(X, Y, S)$, where S is a four bit immediate field. In an example, if one of the bits (e.g., the most significant bit) of the immediate field is zero, the remaining bits select one of the sets of common assignments, and the execution unit routes data from the input vector registers to the multiplication units 502 based on the remaining three bits (thus

replacing the vmulH instruction). If the most significant bit of the immediate field is one, the remaining bits select a state register, which stores assignment values assigning locations in the input vectors to the multiplication units 502.

[0036] Since the state registers are directly specified via an immediate field, one needn't declare variables to assign to register file registers. As such, the values stored in the state registers are persistent and there is no need to re-load them after different execution cycles. This avoids the latency associated with the vmulR instruction described herein. As such, one of the low latency execution units 214 described herein may include a set of state registers storing pre-loaded assignment values specifying locations of input vectors to specific inputs of operations of the low latency execution unit 214. The state registers used in execution of a low latency instruction is identified directly via an immediate field, thereby avoiding latency incurred in assigning registers to declared routing variables.

[0037] In various embodiments, the single sample processing core 112 includes a plurality of low latency execution units 214 for executing instructions used in the performance of asynchronous sampling rate conversion (ASRC) on streams of data received via the audio transport block 104. For example, the single sample processing core 112 employs an infinite impulse response (IIR) ASRC phasor filter including an input stage and an output stage.

[0038] FIG. 6 provides a high level block diagram of an example IIR ASRC real phasor filter employed. The phasor filter includes an input stage and an output stage. As shown, an input sample $X[m_i]$ is received (e.g., directly from the audio fabric block 108, retrieved from system memory, etc.) and summed with a plurality of signals generated by feedback paths contained in the input stage. In some embodiments, there are a plurality of pairs of feedback paths, with each pair including a real path and an imaginary path. The feedback paths include Z^{-1} blocks, which are essentially delay units. The outputs of the delay units are multiplied by respective pole coefficients (denoted in FIG. 6 by $c_n \cdot \cos(\theta_n)$ and by $c_n \cdot \sin(\theta_n)$). The pole coefficient θ_n controls the resonance frequencies of the phasor filter and the pole coefficient c_n controls the magnitude of the pole positions and the decay time of the phasor filter.

[0039] Poles of the phasor filter may be defined by the following relationship:

$$P_n = e^{\text{real}(s_n * T)} * (\cos(\text{imag}(s_n * T)) + i * \sin(\text{imag}(s_n * T))) \quad (1)$$

where s_n corresponds to an N_p s-plane pole of a transfer function of an analog phasor filter and $T = 1/F_s$, where F_s corresponds to an input sampling rate. In other words, the pole coefficients θ_n correspond to the imaginary component of the value $s_n * T$ and the pole coefficients c_n correspond to the exponential of the real component of the value $s_n * T$. Accordingly, to efficiently apply the filter input stage, one must compute the values $c_n * \cos(\theta_n)$ and by $c_n * \sin(\theta_n)$ in advance of the filtering. This involves the computation of the sine and cosine of the imaginary component of $s_n * T$ and the computation of an exponential of the real component of $s_n * T$. For more information regarding design of the phasor filter, see U.S. Patent No. 8,618,961, hereby incorporated by reference in its entirety.

[0040] In the example shown, the ASRC phasor filter also includes an output stage. The output stage includes time-varying output stage time coefficients (denoted in FIG. 3 by $pCurRe(m)$ and $pCurIm(m)$). In various embodiments, the output stage time coefficients are computed via the following relationship:

$$PCur_n^{\Delta t} = c_n^{\Delta t} * (\cos(\theta_n * \Delta t) + \sin(\theta_n * \Delta t)) \quad (2)$$

where Δt represents an elapsed time increment. As such, the output stage time coefficients represent scaled versions of the input stage filter time coefficient values.

[0041] The application of the phasor filter shown in FIG. 6 may be described as follows. The input filter coefficients are applied to an input sample to generate a plurality of filtered sample values. The filtered sample values are then stored in a data buffer (e.g., a register file). The output stage of the filter is then executed a number of times on the filtered sample values depending on an expected time of arrival for another input sample (e.g., based on the input sampling rate). Each time the output stage of the filter is executed, the input filter coefficients are scaled based on an elapsed time increment and the scaled input filter coefficients are applied to corresponding ones of the filtered sample values to generate a plurality of scaled sample values. Each plurality of scaled sample values is then summed together to generate an output sample. Thus, by varying the number of times that the output

stage is executed per input sample, the relationship between the output and input sampling rates of the ASRC algorithm is varied.

[0042] To perform the ASRC algorithm, it is necessary to compute various outputs from the magnitude and angle coefficients of the phasor filter used. In this regard, the single sample processing core 112 described herein with respect to FIGS. 1-2 includes an execution unit (or units) for computing these outputs. FIG. 7 shows high level diagram of an execution sequence 700 for computing these various outputs. As shown, the execution sequence 700 involves the execution of three separate instructions: an exponential instruction, a cosine compute instruction, and a sine compute instruction. Each of these instructions takes a single input (e.g., c_n , θ_n) and generates a single output. In an example, the execution sequence 700 is executed via one of the base execution units 216 including, for example, an exponential compute unit 704, a cosine compute unit 710, and a sine compute unit 716. Assuming that the base pipelines 212 enable two execution cycles per instruction, the execution pipeline is described as follows:

Cycle	Fetch	Decode	Execute 0	Execute 1	Register Write
0	exp				
1	cos	exp			
2	sin	cos	exp		
3		sin	cos	exp	
4			sin	cos	$\exp(\alpha)$
5				sin	$\cos(\theta)$
6					$\sin(\theta)$

[0043] Thus, the computation of the exponential of, cosine, and sine incurs 7 cycles of latency and three cycles of throughput.

[0044] In some embodiments, the single sample processing core 112 includes a low latency execution unit 214 specifically designed to reduce the latency in performance of these

calculations over the execution sequence 700 described with respect to FIG. 7. FIG. 8 provides a schematic illustration of a datapath 800 of such an execution unit. As shown, the datapath 800 includes a first vector register 802, a low latency execution unit 214, and a second vector register 804. In some embodiments, the first and second vector registers 802 and 804 are both components of the same vector register file. As shown, the first and second vector registers 802 and 804 each store four element vectors, but they may store vectors having any number of elements greater than two. The low latency execution unit 214 includes an exponential compute unit, a cosine compute unit, and a sine compute unit. The exponential compute unit, a cosine compute unit, and a sine compute unit may compute the outputs using a number of methods including piecewise polynomial approximations, Newton-Raphson iteration, or other methods.

[0045] In various embodiments, the low latency execution unit 214 is configured to execute a single program instruction taking a two inputs (e.g., a logarithm of c_n and θ_n) to compute three outputs (e.g., an exponential of the logarithm of c_n , $\cos(\theta_n)$, and $\sin(\theta_n)$). For example, this instruction may have the syntax $Y = \text{CexpC}(x)$. An execution pipeline for this instruction may be as depicted in the following table.

Table 1

Cycle	Fetch	Decode	Execute 0	Execute 1	Register Write
0	CexpC				
1		CexpC			
2			CexpC		
3				CexpC	
4					$\sin(\theta) \mid \cos(\theta) \mid \exp(\alpha)$

[0046] Thus, the low latency execution unit 214 only incurs 5 cycles latency and one cycle throughput. This is three times faster than the execution sequence 700 described with respect to FIG. 7, thus providing a significant reduction in latency in the performance of the ASRC algorithm. In some embodiments, such a low latency execution unit 214 may be used to generate coefficient values for polynomial approximations for various functions.

[0047] In various embodiments, a ratio of an input sampling rate of a stream input to the single sample processing core 112 to an output sampling rate generated via the ASRC filter described herein differs with time. Thus, when up-sampling, the number of times that the output stage of the ASRC filter must be executed per input sample received is not known in advance. Accordingly, some instruction sets implementing the ASRC filter described herein employ conditional branching where the output stage of the ASRC filter is executed until a time condition is met (e.g., a period from the previous input sample reaches a predetermined value). However, conditional testing and branching adds extra latency to this filtering.

[0048] Accordingly, in some embodiments, the single sample processing core 112 includes an additional low latency execution unit 214 for executing an instruction that performs fixed-point division to generate a quotient of an elapsed time in units of input samples (“X”) and an output sample time increment in units of input samples (“Y”). The value X is a fractional value and is incremented depending on a ratio of the output sampling rate to the input sampling rate. In some embodiments, each time an input sample is received, a new value of X is used to re-compute the quotient using an updated value of Y to determine a number of times needed to execute the output stage of the ASRC filter. Such a computation enables the use of zero overhead loops in terms of determining how many times to execute the output filter stage when up-sampling, and thus avoids the extra latency incurred from conditional branching.

[0049] FIG. 9 is a block diagram of a various operations performed via the low latency execution unit 214 to compute the quotient. As shown, a first register 902 stores a value for X and a second register 904 stores a value for Y. In the example shown, each value of X and Y is a sixty four bit value, with the first 38 bits representing an integer part and the remaining 26 bits representing the fractional part. In some embodiments, a count of the number of input samples associated with a stream received via the multi-core audio processor 102 is received from the audio fabric block 108, which includes a number of sample counters configured to increment each time samples are stored in a corresponding one of the transport buffers 106. The count of the number of samples is then used to compute values for X and Y.

[0050] As shown, a selection bit routes the fractional part of X to a subtraction unit 906, where a predetermined value 908 is subtracted from the fractional part of X to form a dividend 912 for the computation of the quotient. In various embodiments, the predetermined value is an inverse of an input sampling rate (i.e., a sampling period) associated with the input stream. As such, the output of the subtraction unit 906 represents a fractional time interval remaining until the next input sample is received. Additionally, a control bit determines if any of the thirty eight most significant bits of Y (e.g., representing an integer part of Y) are non-zero. If so, it is indicative that an interval between output samples is too long (i.e., because, when up-sampling, there should be multiple output samples per input sample). In such a case, the execution unit waits until receipt of the next input sample to re-calculate the quotient. In this regard, if any of the thirty eight most significant bits of Y are non-zero, the execution unit advances to a selection block 916 where a select bit is enabled to cause a delay in the execution of the quotient until receipt of the next input sample.

[0051] However, if none of the thirty eight most significant bits of Y are non-zero, a portion of Y (e.g., the 27 most significant bits of Y) form a divisor 910 for computation of the quotient. In the example shown, a Sweeney, Robertson, and Tocher (SRT) division algorithm is employed to compute a six bit quotient. While a standard divider could be used to compute the quotient, the use of an SRT division algorithm enables the computation to be performed with a modest amount of hardware.

[0052] The quotient is then provided to a selection block 916 where the quotient is compared to a predetermined value. In various embodiments, the selection block 916 selects the minimum between the integer above the quotient (i.e., $\text{ciel}(\text{quotient})$) and a predetermined value. For example, if a six bit quotient is used, the predetermined value may correspond to sixty three to support an up-sampling factor of up to two hundred and fifty five.

[0053] The output of the selection block 916 represents a number of output samples generated per input sample processed (e.g., via the input stage of the ASRC filter described herein). As such, the predetermined value essentially sets a cap on the up-sampling factor that the ASRC filter can be used to produce. The cap may be restricted to facilitate the

stability of the ASRC algorithm. Thus, each time an input sample is received, the operations shown in FIG. 9 may be performed to determine in advance a number of times to execute the output filter stage. This avoids the extra latency incurred when conditional testing and branching is included in the filter output stage instruction sets.

[0054] The flowchart and block diagrams in various Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present technology. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

WHAT IS CLAIMED IS:

1. A multi-core audio processor comprising:
 - a data protocol interface configured to receive a stream of audio data;
 - a plurality of data processing cores including a single sample processing core and a block data processing core; and
 - an audio fabric block configured to route samples of the stream between the data protocol interface and the plurality of data processing cores,
 - wherein the single sample processing core comprises an execution unit configured to execute one or more low latency instructions for performing computations for the samples.
2. The multi-core audio processor of claim 1, wherein the execution unit is a first vector execution unit comprising a cosine operation unit, a sine operation unit and an exponential operation unit, the first vector execution unit configured to compute first, second and third outputs in a single execution cycle of the first low latency instruction.
3. The multi-core audio processor of claim 1, wherein the single sample processing core comprises:
 - an instruction memory interface communicably coupled to a memory, the memory storing a first low latency instruction, the first low latency instruction taking at least two inputs;
 - a data memory interface communicably coupled to the audio fabric block; and
 - a first vector register file communicably coupled to the data memory interface, the vector register file including a first vector register storing the at least two inputs, wherein the execution unit includes a first vector execution unit communicably coupled to the first vector register file, wherein the first vector execution unit is configured to perform multiple operations on one of the at least two inputs in a single execution of the low latency instruction and store the results in a second vector register of the vector register file.
4. The multi-core audio processor of claim 2, wherein the single sample processing core further comprises a second vector execution unit, the second vector execution

unit including a multiplication operation unit configured to multiply the third output by a combination of the first and second outputs of the first vector execution unit to generate a filter coefficient.

5. The multi-core audio processor of claim 1, wherein the single sample processing core further comprises an additional execution unit configured to generate a filtered sample based on input sampling rate coefficients applied to an input sample received from the audio fabric block and to store the filtered sample in a data buffer.

6. The multi-core audio processor of claim 5, wherein the additional execution unit is configured to apply output sampling rate coefficients to generate a set of scaled samples used to generate output samples at a different sampling rate than a sampling rate of the stream.

7. The multi-core audio processor of claim 1, wherein the single sample processing core further comprises:
a vector register file; and
a vector execution unit communicably coupled to the vector register file, the vector execution unit having a plurality of hard-wired routes coupling the vector register file to multiplication units of the vector execution unit.

8. The multi-core audio processor of claim 7, further comprising a set of state registers communicably coupled to the vector execution unit, the state registers storing a plurality of assignment vectors assigning vector locations in the vector register file to the multiplication units, the vector execution unit configurable to selectively retrieve an assignment vector from the set of state registers based on a field input.

9. A multi-core audio processor comprising:
a data protocol interface configured to receive a stream of audio data, the stream of audio data having a first sampling rate;

a plurality of data processing cores including a single sample processing core and a block data processing core;

an audio fabric block configured to route samples of the stream between the data protocol interface and the plurality of data processing cores;

wherein the single sample processing core comprises a first execution unit configured to execute a first instruction for computing multiple values used to filter a sample in a single instruction cycle, the first instruction comprising one or both of a complex exponential instruction and a sample rate conversion instruction.

10. The multi-core audio processor of claim 9, wherein the single sample processing core comprises:

a register file communicably coupled to the first execution unit, the register file configured to store an angle and a magnitude of a pole coefficient of a phasor filter, wherein the first instruction configures the first execution unit to compute a sine of the angle, a cosine of the angle, and an exponential function of the magnitude.

11. The multi-core audio processor of claim 10, wherein the register file is a vector register file storing the angle and the magnitude in a single vector register thereof.

12. The multi-core audio processor of claim 9, wherein the single sample processing core further comprises a second execution unit configured to execute a second instruction that computes a quotient based on an elapsed time in units of input samples and an output sample time increment.

13. The multi-core audio processor of claim 12, wherein the audio fabric block further includes an input sample counter communicably coupled to the single sample processing core, the input sample counter configured to count a number of samples of the physical stream received at the data protocol interface, wherein the single sample processing core is configured to compute the elapsed time based on the number of samples of the physical stream and a presentation time of the input sample.

14. The multi-core audio processor of claim 12, wherein the single sample processing core is configured to execute an output stage of a phasor filter a number of times that corresponds to the quotient.

15. The multi-core audio processor of claim 12, wherein the second execution unit executes a Sweeney, Robertson, and Tocher (SRT) division algorithm to generate a six bit quotient.

16. The multi-core audio processor of claim 9, wherein the single sample processing core further comprises a filter coefficient execution unit configured by an instruction to generate filter coefficients based on outputs received from the first execution unit.

17. A single sample audio processing core comprising:
an instruction memory interface communicably coupled to a memory;
a data memory interface configured to receive a sample of audio input data;
a vector register file communicably coupled to the data memory interface;
an execution unit communicably coupled to the vector register file, the execution unit including a plurality of operation units and hard-wired routes from inputs of the operation units to locations of the vector register file; and
a set of state registers communicably coupled to the execution unit, the set of state registers configured to store assignment vectors assigning locations of the vector register file to the inputs of the operation units, wherein the execution unit is configured to execute an instruction having two vector inputs and an immediate field, the immediate field identifying locations of the vector register file or one of the assignment vectors.

18. The single sample audio processing core of claim 17, wherein the execution unit includes a logic unit configured to enable at least one of the hard-wired routes based on a bit value of the immediate field.

19. The single sample audio processing core of claim 18, wherein the execution unit is configured to retrieve an assignment vector assigning locations of the vector register file to the inputs based on the bit value of the immediate field.

20. The single sample audio processing core of claim 17, wherein the plurality of operation units include multipliers configured to multiply values of the two input vectors identified via the immediate field.




FIG. 1




FIG. 2




FIG. 3

	32 bits	32 bits	32 bits	32 bits	32 bits	32 bits	32 bits	32 bits	
V7	V7[7]	V7[6]	V7[5]	V7[4]	V7[3]	V7[2]	V7[1]	V7[0]	304
V6	V6[7]	V6[6]	V6[5]	V6[4]	V6[3]	V6[2]	V6[1]	V6[0]	
V5	V5[7]	V5[6]	V5[5]	V5[4]	V5[3]	V5[2]	V5[1]	V5[0]	
V4	V4[7]	V4[6]	V4[5]	V4[4]	V4[3]	V4[2]	V4[1]	V4[0]	
V3	V3[7]	V3[6]	V3[5]	V3[4]	V3[3]	V3[2]	V3[1]	V3[0]	
V2	V2[7]	V2[6]	V2[5]	V2[4]	V2[3]	V2[2]	V2[1]	V2[0]	
V1	V1[7]	V1[6]	V1[5]	V1[4]	V1[3]	V1[2]	V1[1]	V1[0]	
V0	V0[7]	V0[6]	V0[5]	V0[4]	V0[3]	V0[2]	V0[1]	V0[0]	

FIG. 4




FIG. 5



FIG. 6




FIG. 7




FIG. 8




FIG. 9

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/052346

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F15/78
ADD.

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)
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 practicable, search terms used)

EP0-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2016/275043 A1 (GROCHOWSKI EDWARD T [US] ET AL) 22 September 2016 (2016-09-22) abstract figures 1, 2A, 2B, 4, 5, 13, 19 paragraph [0032] - paragraph [0045] paragraph [0050] paragraph [0057] - paragraph [0068] paragraph [0080] - paragraph [0081] paragraph [0090] - paragraph [0114] -----	1-20
X	US 2011/249744 A1 (BAILEY NEIL [GB]) 13 October 2011 (2011-10-13) the whole document ----- -/--	1-20

☒ Further documents are listed in the continuation of Box C.

☒ See patent family 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 application or patent 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

12 December 2018

Date of mailing of the international search report

21/12/2018

Name and mailing address of the ISA/

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

Authorized officer

Wirtz, Hanno

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/052346

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	HIROAKI SHIKANO ET AL: "Heterogeneous Multi-Core Architecture That Enables 54x AAC-LC Stereo Encoding", IEEE JOURNAL OF SOLID-STATE CIRCUITS., vol. 43, no. 4, 30 April 2008 (2008-04-30), pages 902-910, XP055533711, PISCATAWAY, NJ, USA ISSN: 0018-9200, DOI: 10.1109/JSSC.2008.917531 the whole document -----	1-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2018/052346

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2016275043	A1	22-09-2016	NONE

US 2011249744	A1	13-10-2011	NONE
