



(19) **United States**

(12) **Patent Application Publication**
Ganapathy et al.

(10) **Pub. No.: US 2006/0112260 A1**

(43) **Pub. Date: May 25, 2006**

(54) **METHOD AND APPARATUS OF INSTRUCTION EXECUTION FOR SIGNAL PROCESSORS**

Publication Classification

(51) **Int. Cl.**
G06F 9/30 (2006.01)

(52) **U.S. Cl.** 712/213

(57) **ABSTRACT**

An instruction set architecture (ISA) for application specific signal processor (ASSP) is tailored to digital signal processing applications. The instruction set architecture implemented with the ASSP, is adapted to DSP algorithmic structures. The instruction word of the ISA is typically 20 bits but can be expanded to 40-bits to control two instructions to be executed in series or parallel. All DSP instructions of the ISA are dyadic DSP instructions performing two operations with one instruction in one cycle. The DSP instructions or operations in the preferred embodiment include a multiply instruction (MULT), an addition instruction (ADD), a minimize/maximize instruction (MIN/MAX) also referred to as an extrema instruction, and a no operation instruction (NOP) each having an associated operation code ("opcode"). The present invention efficiently executes DSP instructions by means of the instruction set architecture and the hardware architecture of the application specific signal processor.

(76) Inventors: **Kumar Ganapathy**, Mountain View, CA (US); **Ruban Kanapathipillai**, Fremont, CA (US)

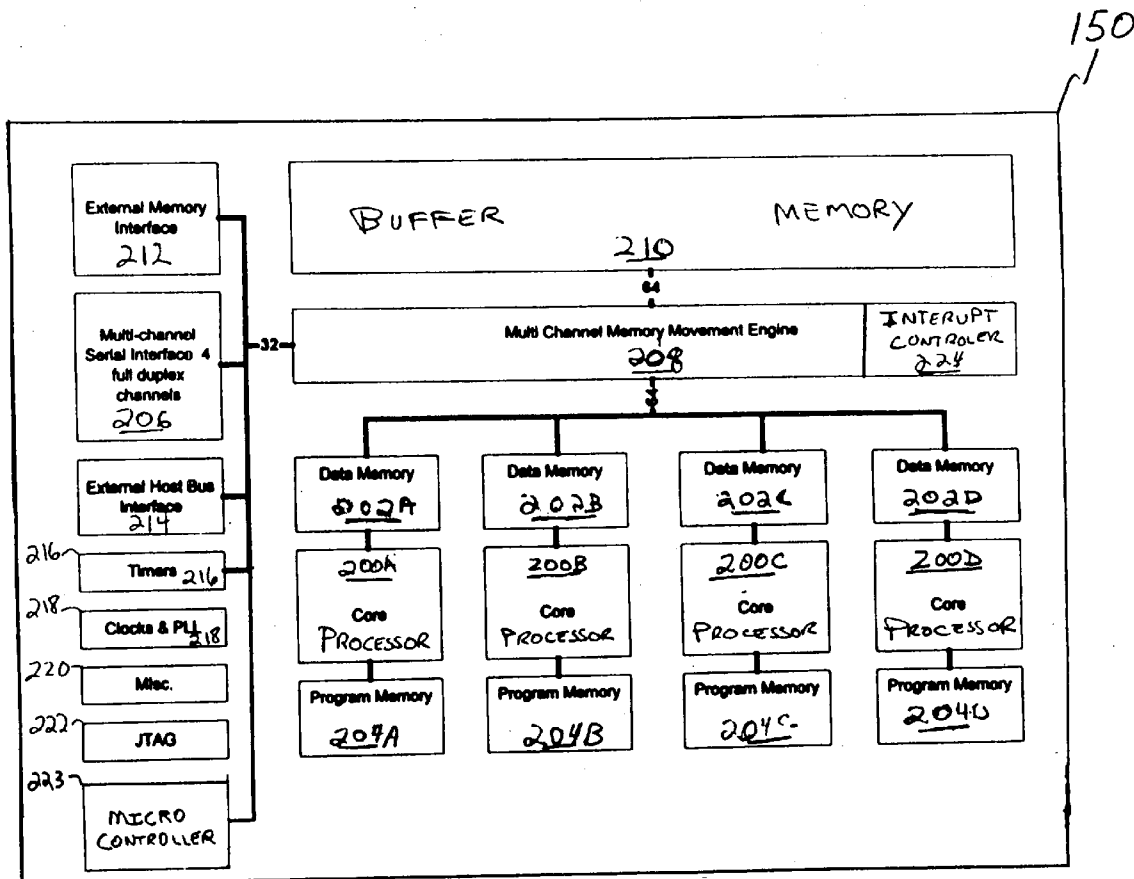
Correspondence Address:
BLAKELY SOKOLOFF TAYLOR & ZAFMAN
12400 WILSHIRE BOULEVARD
SEVENTH FLOOR
LOS ANGELES, CA 90025-1030 (US)

(21) Appl. No.: **11/323,253**

(22) Filed: **Dec. 30, 2005**

Related U.S. Application Data

(62) Division of application No. 10/211,387, filed on Aug. 2, 2002, now Pat. No. 6,988,184, which is a division of application No. 09/494,608, filed on Jan. 31, 2000, now Pat. No. 6,446,195.



100




FIG. 1A




FIG. 1B

150

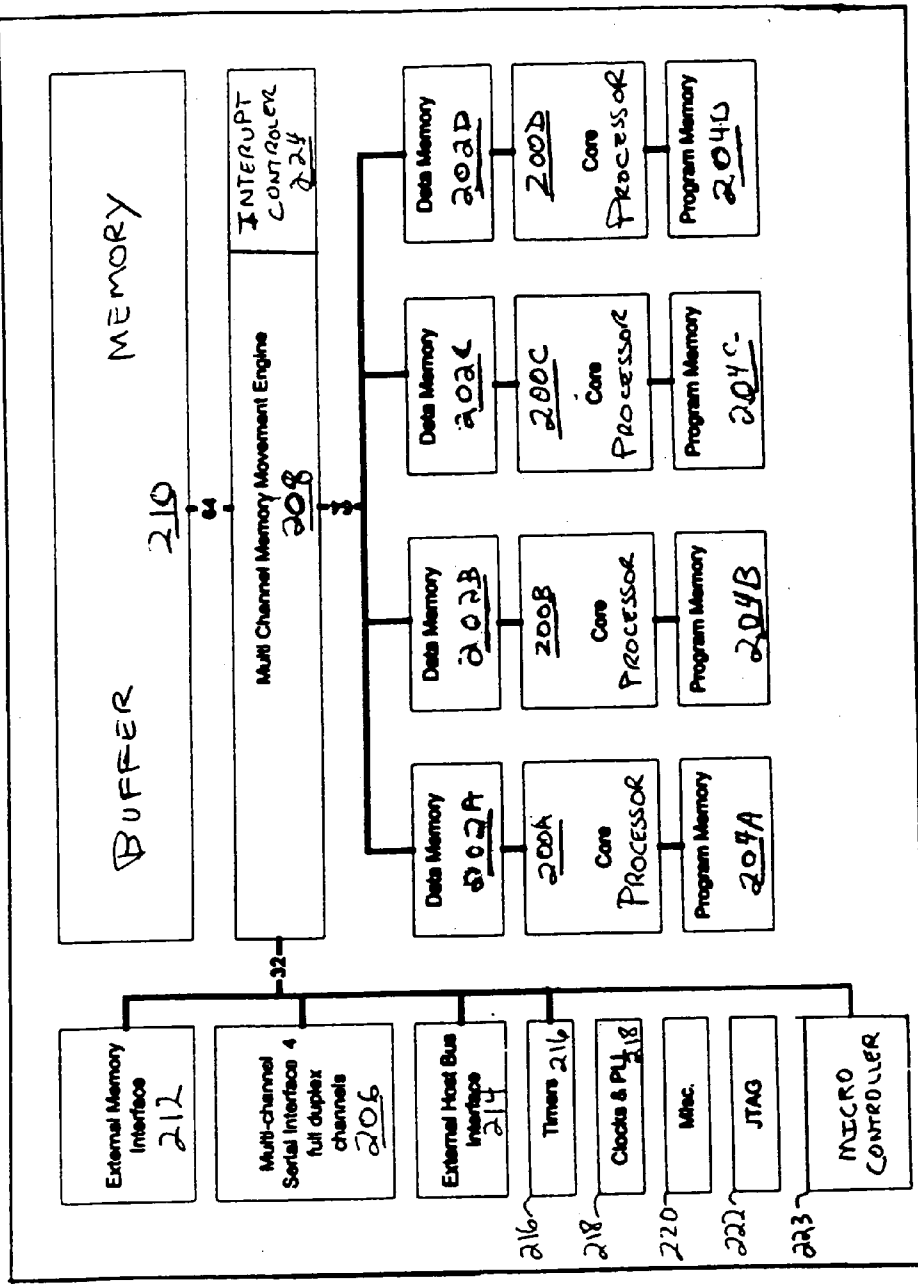


FIG. 2

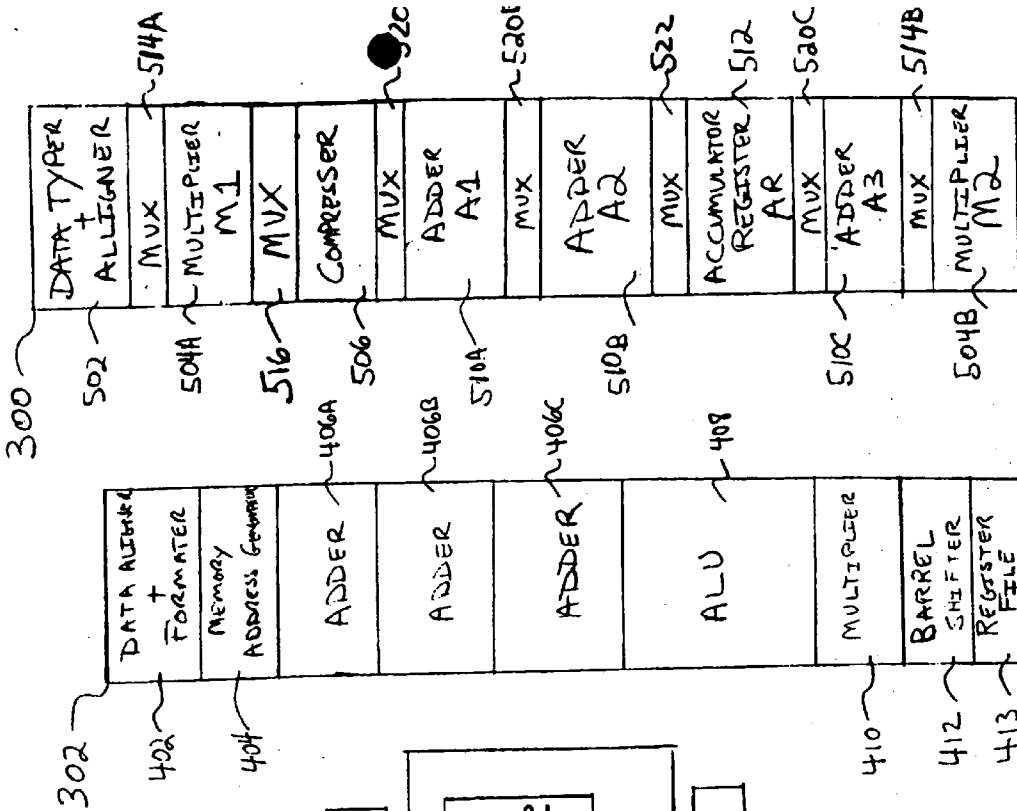


FIG. 5A

FIG. 4

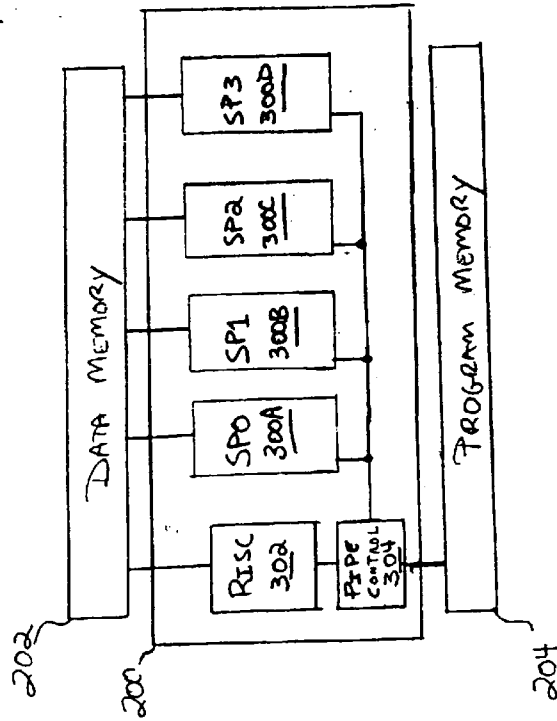


FIG. 3

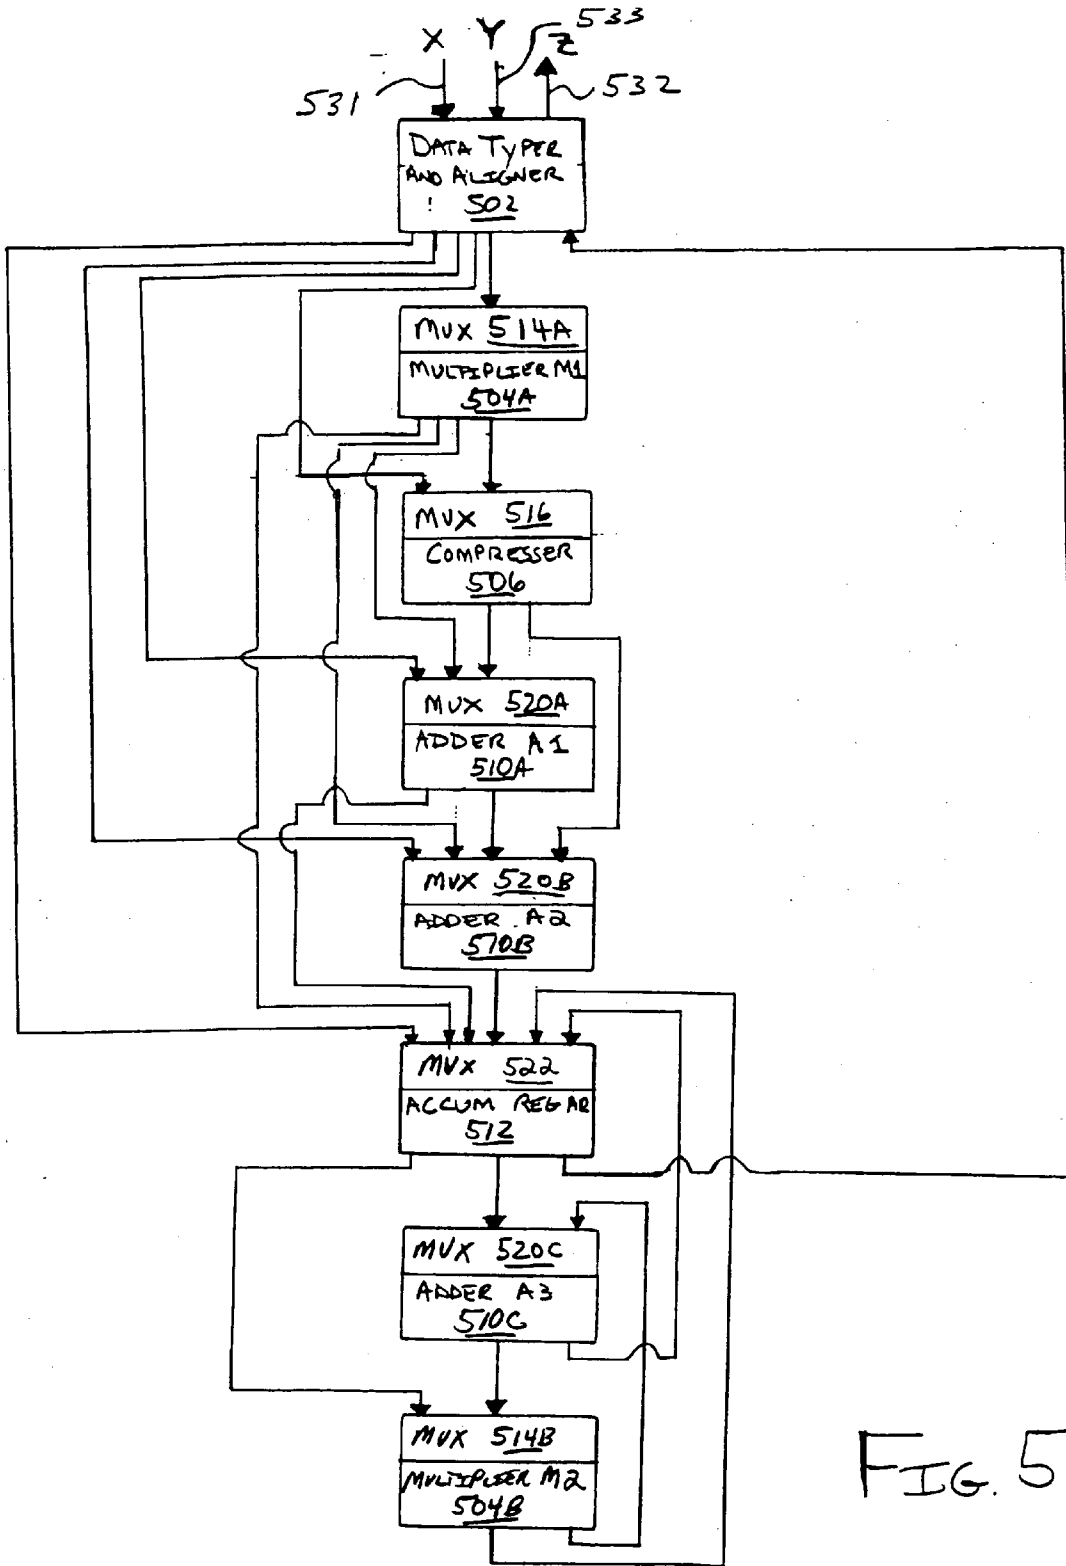


FIG. 5B

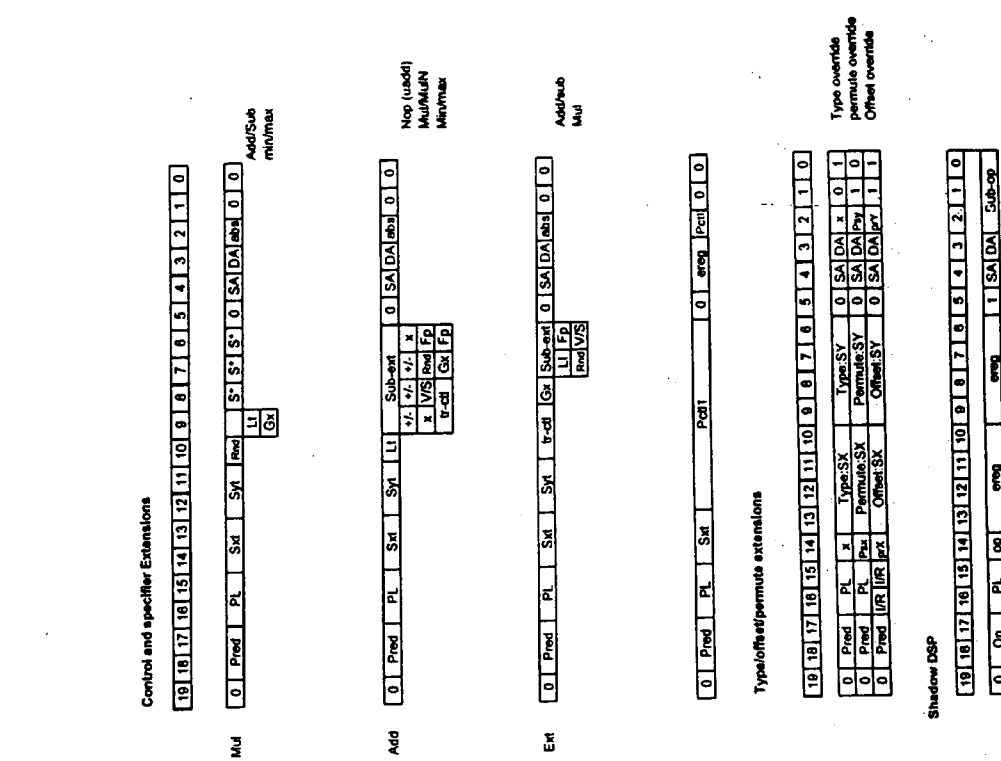
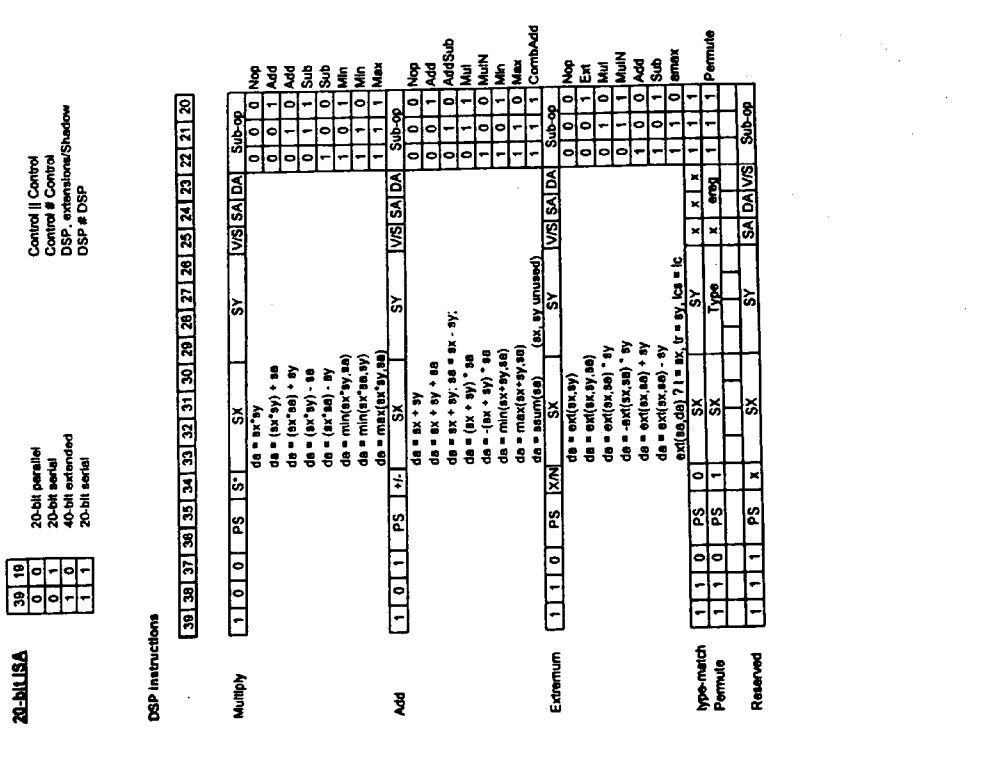


FIG. 6 E



<B11, Bits 9-6> == U15 (Shift Amount)

<B13, Bits 13-10> == U15 .POS

Control Instructions

	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
add_snb	L	Pred	0	0	0	0	RX													+/	
max_min	L	Pred	0	0	0	0	RX														X/M
Shift	L	Pred	0	0	0	0	RX														U11
Logic	L	Pred	0	0	0	0	RX														U11
Mux	L	Pred	0	0	0	0	RX														B, 1, B1
mov	L	Pred	0	0	0	0	RX														Pd
addi	L	Pred	0	0	0	0	RX														
mov2reg	L	Pred	0	0	0	0	S14														
Ldm	L	Pred	0	0	0	0	RX														
Set4bits	L	Pred	0	0	0	0	U14:POS														
Set2bits	L	Pred	0	0	0	0	U14:POS														
Setbit	L	Pred	0	0	0	0	U14:POS														
Move	L	Pred	0	0	0	0	S18														
Jump	L	Pred	0	0	0	0	S19														
Call	L	Pred	0	0	0	0	S19														
Loop	L	Pred	0	0	0	0	U15: Lcount														
Jump	L	Pred	0	0	0	0	RX														
Call	L	Pred	0	0	0	0	RX														
Loop	L	Pred	0	0	0	0	RX														
Test	L	Pred	0	0	0	0	RX														
Testbit	L	Pred	0	0	0	0	RX														
Andp_orp	L	Pred	0	0	0	0	Pe														
Load	L	Pred	0	0	0	0	MX														
Store	L	Pred	0	0	0	0	MZ														
eLoad	L	Pred	0	0	0	0	MX														
eStore	L	Pred	0	0	0	0	MZ														
Extended	L	Pred	0	0	0	0	RX														
Logic2	L	Pred	0	0	0	0	RX														
mov_arg	L	Pred	0	0	0	0	unit														
Cdb	L	Pred	0	0	0	0	RX														
Parity	L	Pred	0	0	0	0	RX														
Str	L	Pred	0	0	0	0	MZ														
Abs	L	Pred	0	0	0	0	RX														
Neg	L	Pred	0	0	0	0	RX														
Div-step	L	Pred	0	0	0	0	RX														
Test & Set	L	Pred	0	0	0	0	RX														
Return	L	Pred	0	0	0	0	unit														
Zero-ec	L	Pred	0	0	0	0	RZ														
eSync	L	Pred	0	0	0	0	RZ														
Swf	L	Pred	0	0	0	0	U13														
Nop	L	Pred	0	0	0	0	U13														

FIG. 6 F

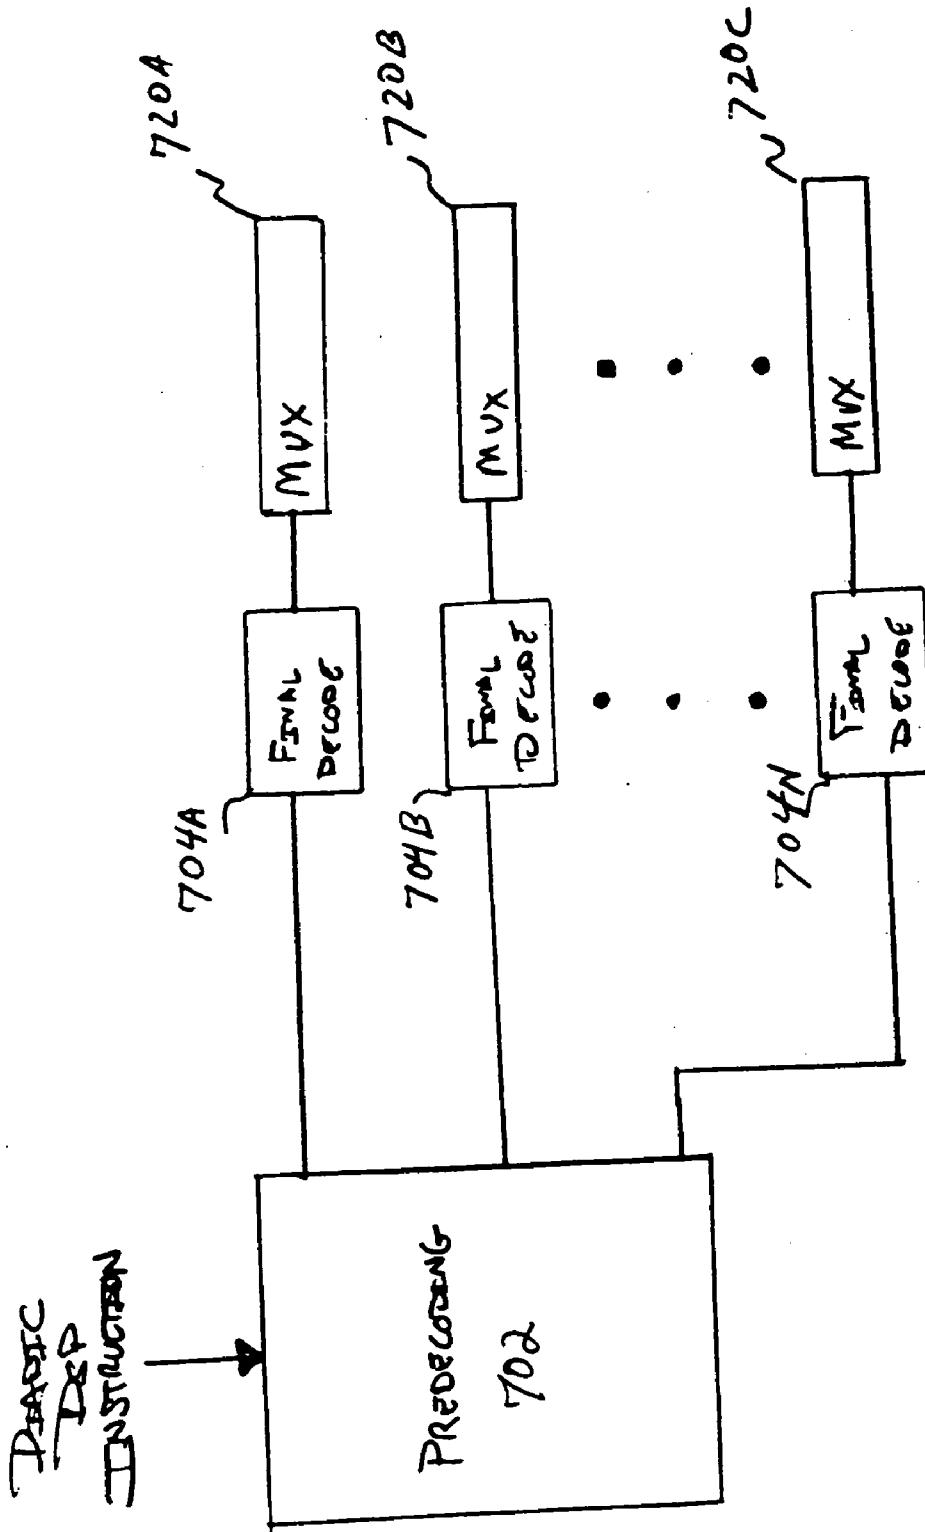


FIG. 7

METHOD AND APPARATUS OF INSTRUCTION EXECUTION FOR SIGNAL PROCESSORS

FIELD OF THE INVENTION

[0001] This invention relates generally to the instruction set architectures (ISA) of processors. More particularly, the invention relates to instruction set architectures for the execution of operations within a signal processing integrated circuit.

BACKGROUND OF THE INVENTION

[0002] Single chip digital signal processing devices (DSP) are relatively well known. DSPs generally are distinguished from general purpose microprocessors in that DSPs typically support accelerated arithmetic operations by including a dedicated multiplier and accumulator (MAC) for performing multiplication of digital numbers. The instruction set for a typical DSP device usually includes a MAC instruction for performing multiplication of new operands and addition with a prior accumulated value stored within an accumulator register. A MAC instruction is typically the only instruction provided in prior art digital signal processors where two DSP operations, multiply followed by add, are performed by the execution of one instruction. However, when performing signal processing functions on data it is often desirable to perform other DSP operations in varying combinations.

[0003] An area where DSPs may be utilized is in telecommunication systems. One use of DSPs in telecommunication systems is digital filtering. In this case a DSP is typically programmed with instructions to implement some filter function in the digital or time domain. The mathematical algorithm for a typical finite impulse response (FIR) filter may look like the equation $Y_n = h_0 X_n + h_1 X_{n-1} + h_2 X_{n-2} + \dots + h_N X_{n-N}$ where h_n are fixed filter coefficients numbering from 1 to N and X_n are the data samples. The equation Y_n may be evaluated by using a software program. However in some applications, it is necessary that the equation be evaluated as fast as possible. One way to do this is to perform the computations using hardware components such as a DSP device programmed to compute the equation Y_n . In order to further speed the process, it is desirable to vectorize the equation and distribute the computation amongst multiple DSPs such that the final result is obtained more quickly. The multiple DSPs operate in parallel to speed the computation process. In this case, the multiplication of terms is spread across the multipliers of the DSPs equally for simultaneous computations of terms. The adding of terms is similarly spread equally across the adders of the DSPs for simultaneous computations. In vectorized processing, the order of processing terms is unimportant since the combination is associative. If the processing order of the terms is altered, it has no effect on the final result expected in a vectorized processing of a function.

[0004] In typical micro processors, a MAC operation would require a multiply instruction and an add instruction to perform both multiplication and addition. To perform these two instructions would require two processing cycles. Additionally, a program written for the typical micro processor would require a larger program memory in order to store the extra instructions necessary to perform the MAC operation. In prior art DSP devices, if a DSP operation other than a MAC DSP instruction need be performed, the opera-

tion requires separate arithmetic instructions programmed into program memory. These separate arithmetic instructions in prior art DSPs similarly require increased program memory space and processing cycles to perform the operation when compared to a single MAC instruction. It is desirable to reduce the number of processing cycles when performing DSP operations. It is desirable to reduce program memory requirements as well.

[0005] DSPs are often programmed in a loop to continuously perform accelerated arithmetic functions including a MAC instruction using different operands. Often times, multiple arithmetic instructions are programmed in a loop to operate on the same data set. The same arithmetic instruction is often executed over and over in a loop using different operands. Additionally, each time one instruction is completed, another instruction is fetched from the program stored in memory during a fetch cycle. Fetch cycles require one or more cycle times to access a memory before instruction execution occurs. Because circuits change state during a fetch cycle, power is consumed and thus it is desirable to reduce the number of fetch cycles. Typically, approximately twenty percent of power consumption may be utilized in the set up and clean up operations of a loop in order to execute DSP instructions. Typically, the loop execution where signal processing of data is performed consumes approximately eighty percent of power consumption with a significant portion being due to instruction fetching. Additionally, because data sets that a DSP device process are usually large, it is also desirable to speed instruction execution by avoiding frequent fetch cycles to memory.

[0006] Additionally, the quality of service over a telephone system often relates to the processing speed of signals. That is particularly the case when a DSP is to provide voice processing, such as voice compression, voice decompression, and echo cancellation for multiple channels. More recently, processing speed has become even more important because of the desire to transmit voice aggregated with data in a packetized form for communication over packetized networks. Delays in processing the packetized voice signal tend to result in the degradation of signal quality on receiving ends.

[0007] It is desirable to provide improved processing of voice and data signals to enhance the quality of voice and data communication over packetized networks. It is desirable to improve the efficiency of using computing resources when performing signal processing functions.

BRIEF SUMMARY OF THE INVENTION

[0008] Briefly, the present invention includes a method, apparatus and system as described in the claims. Multiple application specific signal processor (ASSP) having the instruction set architecture of the present invention, including the dyadic DSP instructions, are provided within gateways in communication systems to provide improved voice and data communication over a packetized network. Each ASSP includes a serial interface, a buffer memory, and four core processors for each to simultaneously process multiple channels of voice or data. Each core processor preferably includes a reduced instruction set computer (RISC) processor and four signal processing units (SPs). Each SP includes multiple arithmetic blocks to simultaneously process multiple voice and data communication signal samples for

communication over IP, ATM, Frame Relay or other packetized network. The four signal processing units can execute the digital signal processing algorithms in parallel. Each ASSP is flexible and can be programmed to perform many network functions or data/voice processing functions, including voice and data compression/decompression in telecommunications systems (such as CODECs) particularly packetized telecommunication networks, simply by altering the software program controlling the commands executed by the ASSP.

[0009] An instruction set architecture for the ASSP is tailored to digital signal processing applications including audio and speech processing such as compression/decompression and echo cancellation. The instruction set architecture implemented with the ASSP, is adapted to DSP algorithmic structures. This adaptation of the ISA of the present invention to DSP algorithmic structures balances the ease of implementation, processing efficiency, and programmability of DSP algorithms. The instruction set architecture may be viewed as being two component parts, one (RISC ISA) corresponding to the RISC control unit and another (DSP ISA) to the DSP datapaths of the signal processing units 300. The RISC ISA is a register based architecture including 16-registers within the register file 413, while the DSP ISA is a memory based architecture with efficient digital signal processing instructions. The instruction word for the ASSP is typically 20 bits but can be expanded to 40-bits to control two instructions to be executed in series or parallel, such as two RISC control instructions and extended DSP instructions. The instruction set architecture of the ASSP has four distinct types of instructions to optimize the DSP operational mix. These are (1) a 20-bit DSP instruction that uses mode bits in control registers (i.e. mode registers), (2) a 40-bit DSP instruction having control extensions that can override mode registers, (3) a 20-bit dyadic DSP instruction, and (4) a 40 bit dyadic DSP instruction. These instructions are for accelerating calculations within the core processor of the type where $D=[(A \text{ op1 } B) \text{ op2 } C]$ and each of "op1" and "op2" can be a multiply, add, extremum (min/max) or other primitive DSP class of operation on the three operands A, B, and C. The ISA of the ASSP which accelerates these calculations allows efficient chaining of different combinations of operations. All DSP instructions of the instruction set architecture of the ASSP are dyadic DSP instructions to execute two operations in one instruction with one cycle throughput. A dyadic DSP instruction is a combination of two basic DSP operations in one instruction and includes a main DSP operation (MAIN OP) and a sub DSP operation (SUB OP). Generally, the instruction set architecture of the present invention can be generalized to combining any pair of basic DSP operations to provide very powerful dyadic instruction combinations. The DSP arithmetic instructions or operations in the preferred embodiment include a multiply instruction (MULT), an addition instruction (ADD), a minimize/maximize instruction (MIN/MAX) also referred to as an extrema instruction, and a no operation instruction (NOP) each having an associated operation code ("opcode"). The present invention efficiently executes these dyadic DSP instructions by means of the instruction set architecture and the hardware architecture of the application specific signal processor. For example, the DSP instructions can process vector data or scalar data automatically using a single instruction and provide the appropriate vector or scalar output results.

BRIEF DESCRIPTIONS OF THE DRAWINGS

[0010] FIG. 1A is a block diagram of a system utilizing the present invention.

[0011] FIG. 1B is a block diagram of a printed circuit board utilizing the present invention within the gateways of the system in FIG. 1A.

[0012] FIG. 2 is a block diagram of the Application Specific Signal Processor (ASSP) of the present invention.

[0013] FIG. 3 is a block diagram of an instance of the core processors within the ASSP of the present invention.

[0014] FIG. 4 is a block diagram of the RISC processing unit within the core processors of FIG. 3.

[0015] FIG. 5A is a block diagram of an instance of the signal processing units within the core processors of FIG. 3.

[0016] FIG. 5B is a more detailed block diagram of FIG. 5A illustrating the bus structure of the signal processing unit.

[0017] FIG. 6A is an exemplary instruction sequence illustrating a program model for DSP algorithms employing the instruction set architecture of the present invention.

[0018] FIG. 6B is a chart illustrating the permutations of the dyadic DSP instructions.

[0019] FIG. 6C is an exemplary bitmap for a control extended dyadic DSP instruction.

[0020] FIG. 6D is an exemplary bitmap for a non-extended dyadic DSP instruction.

[0021] FIG. 6E and 6F list the set of 20-bit instructions for the ISA of the present invention.

[0022] FIG. 6G lists the set of extended control instructions for the ISA of the present invention.

[0023] FIG. 6H lists the set of 40-bit DSP instructions for the ISA of the present invention.

[0024] FIG. 6I lists the set of addressing instructions for the ISA of the present invention.

[0025] FIG. 7 is a block diagram illustrating the instruction decoding and configuration of the functional blocks of the signal processing units.

[0026] Like reference numbers and designations in the drawings indicate like elements providing similar functionality. A letter after a reference designator number represents an instance of an element having the reference designator number.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

[0027] In the following detailed description of the present invention, numerous specific details are set forth in order to provide a thorough understanding of the present invention. However, it will be obvious to one skilled in the art that the present invention may be practiced without these specific details. In other instances well known methods, procedures, components, and circuits have not been described in detail so as not to unnecessarily obscure aspects of the present invention. Furthermore, the present invention will be

described in particular embodiments but may be implemented in hardware, software, firmware or a combination thereof

[0028] Multiple application specific signal processors (ASSPs) having the instruction set architecture of the present invention, including dyadic DSP instructions, are provided within gateways in communication systems to provide improved voice and data communication over a packetized network. Each ASSP includes a serial interface, a buffer memory and four core processors in order to simultaneously process multiple channels of voice or data. Each core processor preferably includes a reduced instruction set computer (RISC) processor and four signal processing units (SPs). Each SP includes multiple arithmetic blocks to simultaneously process multiple voice and data communication signal samples for communication over IP, ATM, Frame Relay, or other packetized network. The four signal processing units can execute digital signal processing algorithms in parallel. Each ASSP is flexible and can be programmed to perform many network functions or data/voice processing functions, including voice and data compression/decompression in telecommunication systems (such as CODECs), particularly packetized telecommunication networks, simply by altering the software program controlling the commands executed by the ASSP.

[0029] An instruction set architecture for the ASSP is tailored to digital signal processing applications including audio and speech processing such as compression/decompression and echo cancellation. The instruction set architecture implemented with the ASSP, is adapted to DSP algorithmic structures. This adaptation of the ISA of the present invention to DSP algorithmic structures balances the ease of implementation, processing efficiency, and programmability of DSP algorithms. The instruction set architecture may be viewed as being two component parts, one (RISC ISA) corresponding to the RISC control unit and another (DSP ISA) to the DSP datapaths of the signal processing units 300. The RISC ISA is a register based architecture including 16-registers within the register file 413, while the DSP ISA is a memory based architecture with efficient digital signal processing instructions. The instruction word for the ASSP is typically 20 bits but can be expanded to 40-bits to control two instructions to be executed in series or parallel, such as two RISC control instruction and extended DSP instructions. The instruction set architecture of the ASSP has four distinct types of instructions to optimize the DSP operational mix. These are (1) a 20-bit DSP instruction that uses mode bits in control registers (i.e. mode registers), (2) a 40-bit DSP instruction having control extensions that can override mode registers, (3) a 20-bit dyadic DSP instruction, and (4) a 40 bit dyadic DSP instruction. These instructions are for accelerating calculations within the core processor of the type where $D=[(A \text{ op1 } B) \text{ op2 } C]$ and each of "op1" and "op2" can be a multiply, add or extremum (min/max) class of operation on the three operands A, B, and C. The ISA of the ASSP which accelerates these calculations allows efficient chaining of different combinations of operations.

[0030] All DSP instructions of the instruction set architecture of the ASSP are dyadic DSP instructions to execute two operations in one instruction with one cycle throughput. A dyadic DSP instruction is a combination of two DSP instructions or operations in one instruction and includes a main DSP operation (MAIN OP) and a sub DSP operation

(SUB OP). Generally, the instruction set architecture of the present invention can be generalized to combining any pair of basic DSP operations to provide very powerful dyadic instruction combinations. The DSP arithmetic operations in the preferred embodiment include a multiply instruction (MULT), an addition instruction (ADD), a minimize/maximize instruction (MIN/MAX) also referred to as an extrema instruction, and a no operation instruction (NOP) each having an associated operation code ("opcode").

[0031] The present invention efficiently executes these dyadic DSP instructions by means of the instruction set architecture and the hardware architecture of the application specific signal processor.

[0032] Referring now to FIG. 1A, a voice and data communication system 100 is illustrated. The system 100 includes a network 101 which is a packetized or packet-switched network, such as IP, ATM, or frame relay. The network 101 allows the communication of voice/speech and data between endpoints in the system 100, using packets. Data may be of any type including audio, video, email, and other generic forms of data. At each end of the system 100, the voice or data requires packetization when transceived across the network 101. The system 100 includes gateways 104A, 104B, and 104C in order to packetize the information received for transmission across the network 101. A gateway is a device for connecting multiple networks and devices that use different protocols. Voice and data information may be provided to a gateway 104 from a number of different sources in a variety of digital formats. In system 100, analog voice signals are transceived by a telephone 108. In system 100, digital voice signals are transceived at public branch exchanges (PBX) 112A and 112B which are coupled to multiple telephones, fax machines, or data modems. Digital voice signals are transceived between PBX 112A and PBX 112B with gateways 104A and 104C, respectively. Digital data signals may also be transceived directly between a digital modem 114 and a gateway 104A. Digital modem 114 may be a Digital Subscriber Line (DSL) modem or a cable modem. Data signals may also be coupled into system 100 by a wireless communication system by means of a mobile unit 118 transceiving digital signals or analog signals wirelessly to a base station 116. Base station 116 converts analog signals into digital signals or directly passes the digital signals to gateway 104B. Data may be transceived by means of modem signals over the plain old telephone system (POTS) 107B using a modem 110. Modem signals communicated over POTS 107B are traditionally analog in nature and are coupled into a switch 106B of the public switched telephone network (PSTN). At the switch 106B, analog signals from the POTS 107B are digitized and transceived to the gateway 104B by time division multiplexing (TDM) with each time slot representing a channel and one DS0 input to gateway 104B. At each of the gateways 104A, 104B and 104C, incoming signals are packetized for transmission across the network 101. Signals received by the gateways 104A, 104B and 104C from the network 101 are depacketized and transcoded for distribution to the appropriate destination.

[0033] Referring now to FIG. 1B, a network interface card (NIC) 130 of a gateway 104 is illustrated. The NIC 130 includes one or more application-specific signal processors (ASSPs) 150A-150N. The number of ASSPs within a gateway is expandable to handle additional channels. Line

interface devices 131 of NIC 130 provide interfaces to various devices connected to the gateway, including the network 101. In interfacing to the network 101, the line interface devices packetize data for transmission out on the network 101 and depacketize data which is to be received by the ASSP devices. Line interface devices 131 process information received by the gateway on the receive bus 134 and provides it to the ASSP devices. Information from the ASSP devices 150 is communicated on the transmit bus 132 for transmission out of the gateway. A traditional line interface device is a multi-channel serial interface or a UTOPIA device. The NIC 130 couples to a gateway backplane/network interface bus 136 within the gateway 104. Bridge logic 138 transceives information between bus 136 and NIC 130. Bridge logic 138 transceives signals between the NIC 130 and the backplane/network interface bus 136 onto the host bus 139 for communication to either one or more of the ASSP devices 150A-150N, a host processor 140, or a host memory 142. Optionally coupled to each of the one or more ASSP devices 150A through 150N (generally referred to as ASSP 150) are optional local memory 145A through 145N (generally referred to as optional local memory 145), respectively. Digital data on the receive bus 134 and transmit bus 132 is preferably communicated in bit wide fashion. While internal memory within each ASSP may be sufficiently large to be used as a scratchpad memory, optional local memory 145 may be used by each of the ASSPs 150 if additional memory space is necessary.

[0034] Each of the ASSPs 150 provide signal processing capability for the gateway. The type of signal processing provided is flexible because each ASSP may execute differing signal processing programs. Typical signal processing and related voice packetization functions for an ASSP include (a) echo cancellation; (b) video, audio, and voice/speech compression/decompression (voice/speech coding and decoding); (c) delay handling (packets, frames); (d) loss handling; (e) connectivity (LAN and WAN); (f) security (encryption/decryption); (g) telephone connectivity; (h) protocol processing (reservation and transport protocols, RSVP, TCP/IP, RTP, UDP for IP, and AAL2, AAL1, AAL5 for ATM); (i) filtering; (j) Silence suppression; (k) length handling (frames, packets); and other digital signal processing functions associated with the communication of voice and data over a communication system. Each ASSP 150 can perform other functions in order to transmit voice and data to the various endpoints of the system 100 within a packet data stream over a packetized network.

[0035] Referring now to FIG. 2, a block diagram of the ASSP 150 is illustrated. At the heart of the ASSP 150 are four core processors 200A-200D. Each of the core processors 200A-200D is respectively coupled to a data memory 202A-202D and a program memory 204A-204D. Each of the core processors 200A-200D communicates with outside channels through the multi-channel serial interface 206, the multi-channel memory movement engine 208, buffer memory 210, and data memory 202A-202D. The ASSP 150 further includes an external memory interface 212 to couple to the external optional local memory 145. The ASSP 150 includes an external host interface 214 for interfacing to the external host processor 140 of FIG. 1B.—Further included within the ASSP 150 are timers 216, clock generators and a phase-lock loop 218, miscellaneous control logic 220, and a Joint Test Action Group (JTAG) test access port 222 for boundary scan testing. The multi-channel serial interface

206 may be replaced with a UTOPIA parallel interface for some applications such as ATM. The ASSP 150 further includes a microcontroller 223 to perform process scheduling for the core processors 200A-200D and the coordination of the data movement within the ASSP as well as an interrupt controller 224 to assist in interrupt handling and the control of the ASSP 150.

[0036] Referring now to FIG. 3, a block diagram of the core processor 200 is illustrated coupled to its respective data memory 202 and program memory 204. Core processor 200 is the block diagram for each of the core processors 200A-200D. Data memory 202 and program memory 204 refers to a respective instance of data memory 202A-202D and program memory 204A-204D, respectively. The core processor 200 includes four signal processing units SP0300A, SP1300B, SP2300C and SP3300D. The core processor 200 further includes a reduced instruction set computer (RISC) control unit 302 and a pipeline control unit 304. The signal processing units 300A-300D perform the signal processing tasks on data while the RISC control unit 302 and the pipeline control unit 304 perform control tasks related to the signal processing function performed by the SPs 300A-300D. The control provided by the RISC control unit 302 is coupled with the SPs 300A-300D at the pipeline level to yield a tightly integrated core processor 200 that keeps the utilization of the signal processing units 300 at a very high level.

[0037] The signal processing tasks are performed on the datapaths within the signal processing units 300A-300D. The nature of the DSP algorithms are such that they are inherently vector operations on streams of data, that have minimal temporal locality (data reuse). Hence, a data cache with demand paging is not used because it would not function well and would degrade operational performance. Therefore, the signal processing units 300A-300D are allowed to access vector elements (the operands) directly from data memory 202 without the overhead of issuing a number of load and store instructions into memory resulting, in very efficient data processing. Thus, the instruction set architecture of the present invention having a 20 bit instruction word which can be expanded to a 40 bit instruction word, achieves better efficiencies than VLIW architectures using 256-bits or higher instruction widths by adapting the ISA to DSP algorithmic structures. The adapted ISA leads to very compact and low-power hardware that can scale to higher computational requirements. The operands that the ASSP can accommodate are varied in data type and data size. The data type may be real or complex, an integer value or a fractional value, with vectors having multiple elements of different sizes. The data size in the preferred embodiment is 64 bits but larger data sizes can be accommodated with proper instruction coding.

[0038] Referring now to FIG. 4, a detailed block diagram of the RISC control unit 302 is illustrated. RISC control unit 302 includes a data aligner and formatter 402, a memory address generator 404, three adders 406A-406C, an arithmetic logic unit (ALU) 408, a multiplier 410, a barrel shifter 412, and a register file 413. The register file 413 points to a starting memory location from which memory address generator 404 can generate addresses into data memory 202. The RISC control unit 302 is responsible for supplying addresses to data memory so that the proper data stream is fed to the signal processing units 300A-300D. The RISC

control unit **302** is a register to register organization with load and store instructions to move data to and from data memory **202**. Data memory addressing is performed by RISC control unit using a 32-bit register as a pointer that specifies the address, post-modification offset, and type and permute fields. The type field allows a variety of natural DSP data to be supported as a “first class citizen” in the architecture. For instance, the complex type allows direct operations on complex data stored in memory removing a number of bookkeeping instructions. This is useful in supporting QAM demodulators in data modems very efficiently.

[0039] Referring now to **FIG. 5A**, a block diagram of a signal processing unit **300** is illustrated which represents an instance of the SPs **300A-300D**. Each of the signal processing units **300** includes a data typer and aligner **502**, a first multiplier **M1504A**, a compressor **506**, a first adder **A1510A**, a second adder **A2510B**, an accumulator register **512**, a third adder **A3510C**, and a second multiplier **M2504B**. Adders **510A-510C** are similar in structure and are generally referred to as adder **510**. Multipliers **504A** and **504B** are similar in structure and generally referred to as multiplier **504**. Each of the multipliers **504A** and **504B** have a multiplexer **514A** and **514B** respectively at its input stage to multiplex different inputs from different busses into the multipliers. Each of the adders **510A**, **510B**, **510C** also have a multiplexer **520A**, **520B**, and **520C** respectively at its input stage to multiplex different inputs from different busses into the adders. These multiplexers and other control logic allow the adders, multipliers and other components within the signal processing units **300A-300C** to be flexibly interconnected by proper selection of multiplexers. In the preferred embodiment, multiplier **M1504A**, compressor **506**, adder **A1510A**, adder **A2510B** and accumulator **512** can receive inputs directly from external data buses through the data typer and aligner **502**. In the preferred embodiment, adder **510C** and multiplier **M2504B** receive inputs from the accumulator **512** or the outputs from the execution units multiplier **M1504A**, compressor **506**, adder **A1510A**, and adder **A2510B**.

[0040] Program memory **204** couples to the pipe control **304** which includes an instruction buffer that acts as a local loop cache. The instruction buffer in the preferred embodiment has the capability of holding four instructions. The instruction buffer of the pipe control **304** reduces the power consumed in accessing the main memories to fetch instructions during the execution of program loops.

[0041] Referring now to **FIG. 5B**, a more detailed block diagram of the functional blocks and the bus structure of the signal processing unit is illustrated. Dyadic DSP instructions are possible because of the structure and functionality provided in each signal processing unit. Output signals are coupled out of the signal processor **300** on the Z output bus **532** through the data typer and aligner **502**. Input signals are coupled into the signal processor **300** on the X input bus **531** and Y input bus **533** through the data typer and aligner **502**. Internally, the data typer and aligner **502** has a different data bus to couple to each of multiplier **M1504A**, compressor **506**, adder **A1510A**, adder **A2510B**, and accumulator register **AR 512**. While the data typer and aligner **502** could have data busses coupling to the adder **A3510C** and the multiplier **M2504B**, in the preferred embodiment it does not in order to avoid extra data lines and conserve area usage of an integrated circuit. Output data is coupled from the accu-

mulator register **AR 512** into the data typer and aligner **502**. Multiplier **M1504A** has buses to couple its output into the inputs of the compressor **506**, adder **A1510A**, adder **A2510B**, and the accumulator registers **AR 512**. Compressor **506** has buses to couple its output into the inputs of adder **A1510A** and adder **A2510B**. Adder **A1510A** has a bus to couple its output into the accumulator registers **512**. Adder **A2510B** has buses to couple its output into the accumulator registers **512**. Accumulator registers **512** has buses to couple its output into multiplier **M2504B**, adder **A3510C**, and data typer and aligner **502**. Adder **A3510C** has buses to couple its output into the multiplier **M2504B** and the accumulator registers **512**. Multiplier **M2504B** has buses to couple its output into the inputs of the adder **A3510C** and the accumulator registers **AR 512**.

Instruction Set Architecture

[0042] The instruction set architecture of the ASSP **150** is tailored to digital signal processing applications including audio and speech processing such as compression/decompression and echo cancellation. In essence, the instruction set architecture implemented with the ASSP **150**, is adapted to DSP algorithmic structures. The adaptation of the ISA of the present invention to DSP algorithmic structures is a balance between ease of implementation, processing efficiency, and programmability of DSP algorithms. The ISA of the present invention provides for data movement operations, DSP/arithmetic/logical operations, program control operations (such as function calls/returns, unconditional/conditional jumps and branches), and system operations (such as privilege, interrupt/trap/hazard handling and memory management control).

[0043] Referring now to **FIG. 6A**, an exemplary instruction sequence **600** is illustrated for a DSP algorithm program model employing the instruction set architecture of the present invention. The instruction sequence **600** has an outer loop **601** and an inner loop **602**. Because DSP algorithms tend to perform repetitive computations, instructions **60S** within the inner loop **602** are executed more often than others. Instructions **603** are typically parameter setup code to set the memory pointers, provide for the setup of the outer loop **601**, and other 2×20 control instructions. Instructions **607** are typically context save and function return instructions or other 2×20 control instructions. Instructions **603** and **607** are often considered overhead instructions which are typically infrequently executed. Instructions **604** are typically to provide the setup for the inner loop **602**, other control through 2×20 control instructions, or offset extensions for pointer backup. Instructions **606** typically provide tear down of the inner loop **602**, other control through 2×20 control instructions, and combining of datapath results within the signal processing units. Instructions **605** within the inner loop **602** typically provide inner loop execution of DSP operations, control of the four signal processing units **300** in a single instruction multiple data execution mode, memory access for operands, dyadic DSP operations, and other DSP functionality through the $20/40$ bit DSP instructions of the ISA of the present invention. Because instructions **605** are so often repeated, significant improvement in operational efficiency may be had by providing the DSP instructions, including general dyadic instructions and dyadic DSP instructions, within the ISA of the present invention.

[0044] The instruction set architecture of the ASSP 150 can be viewed as being two component parts, one (RISC ISA) corresponding to the RISC control unit and another (DSP ISA) to the DSP datapaths of the signal processing units 300. The RISC ISA is a register based architecture including sixteen registers within the register file 413, while the DSP ISA is a memory based architecture with efficient digital signal processing instructions. The instruction word for the ASSP is typically 20 bits but can be expanded to 40-bits to control two RISC or DSP instructions to be executed in series or parallel, such as a RISC control instruction executed in parallel with a DSP instruction, or a 40 bit extended RISC or DSP instruction.

[0045] The instruction set architecture of the ASSP 150 has 4 distinct types of instructions to optimize the DSP operational mix. These are (1) a 20-bit DSP instruction that uses mode bits in control registers (i.e. mode registers), (2) a 40-bit DSP instruction having control extensions that can override mode registers, (3) a 20-bit dyadic DSP instruction, and (4) a 40 bit dyadic DSP instruction. These instructions are for accelerating calculations within the core processor 200 of the type where $D=(A \text{ op1 } B) \text{ op2 } C$ and each of "op1" and "op2" can be a multiply, add or extremum (min/max) class of operation on the three operands A, B, and C. The ISA of the ASSP 150 which accelerates these calculations allows efficient chaining of different combinations of operations. Because these type of operations require three operands, they must be available to the processor. However, because the device size places limits on the bus structure, bandwidth is limited to two vector reads and one vector write each cycle into and out of data memory 202. Thus one of the operands, such as B or C, needs to come from another source within the core processor 200. The third operand can be placed into one of the registers of the accumulator 512 or the RISC register file 413. In order to accomplish this within the core processor 200 there are two subclasses of the 20-bit DSP instructions which are (1) A and B specified by a 4-bit specifier, and C and D by a 1-bit specifier and (2) A and C specified by a 4-bit specifier, and B and D by a 1 bit specifier.

[0046] Instructions for the ASSP are always fetched 40-bits at a time from program memory with bit 39 and 19 indicating the type of instruction. After fetching, the instruction is grouped into two sections of 20 bits each for execution of operations. In the case of 20-bit control instructions with parallel execution (bit 39=0, bit 19=0), the two 20-bit sections are control instructions that are executed simultaneously. In the case of 20-bit control instructions for serial execution (bit 39=0, bit 19=1), the two 20-bit sections are control instructions that are executed serially. In the case of 20-bit DSP instructions for serial execution (bit 39=1, bit 19=1), the two 20-bit sections are DSP instructions that are executed serially. In the case of 40-bit DSP instructions (bit 39=1, bit 19=0), the two 20 bit sections form one extended DSP instruction which are executed simultaneously.

[0047] The ISA of the ASSP 150 is fully predicated providing for execution prediction. Within the 20-bit RISC control instruction word and the 40-bit extended DSP instruction word there are 2 bits of each instruction specifying one of four predicate registers within the RISC control unit 302. Depending upon the condition of the predicate register, instruction execution can conditionally change base on its contents.

[0048] In order to access operands within the data memory 202 or registers within the accumulator 512 or register file 413, a 6-bit specifier is used in the DSP extended instructions to access operands in memory and registers. Of the six bit specifier used in the extended DSP instructions, the MSB (Bit 5) indicates whether the access is a memory access or register access. In the preferred embodiment, if Bit 5 is set to logical one, it denotes a memory access for an operand. If Bit 5 is set to a logical zero, it denotes a register access for an operand. If Bit 5 is set to 1, the contents of a specified register (rX where X: 0-7) are used to obtain the effective memory address and post-modify the pointer field by one of two possible offsets specified in one of the specified rX registers. If Bit 5 is set to 0, Bit 4 determines what register set has the contents of the desired operand. If Bit-4 is set to 0, then the remaining specified bits 3:0 control access to the registers within the register file 413 or to registers within the signal processing units 300.

DSP Instructions

[0049] There are four major classes of DSP instructions for the ASSP 150 these are

[0050] 1) Multiply (MULT): Controls the execution of the main multiplier connected to data buses from memory. Controls: Rounding, sign of multiply Operates on vector data specified through type field in address register

[0051] Second operation: Add, Sub, Min, Max in vector or scalar mode

[0052] 2) Add (ADD): Controls the execution of the main-adder Controls: absolute value control of the inputs, limiting the result

[0053] Second operation: Add, add-sub, mult, mac, min, max

[0054] 3) Extremum (MIN/MAX): Controls the execution of the main-adder Controls: absolute value control of the inputs, Global or running max/min with T register, TR register recording control

[0055] Second operation: add, sub, mult, mac, min, max

[0056] 4) Misc: type-match and permute operations.

[0057] The ASSP 150 can execute these DSP arithmetic operations in vector or scalar fashion. In scalar execution, a reduction or combining operation is performed on the vector results to yield a scalar result. It is common in DSP applications to perform scalar operations, which are efficiently performed by the ASSP 150.

[0058] The 20-bit DSP instruction words have 4-bit operand specifiers that can directly access data memory using 8 address registers (r0-r7) within the register file 413 of the RISC control unit 302. The method of addressing by the 20 bit DSP instruction word is regular indirect with the address register specifying the pointer into memory, post-modification value, type of data accessed and permutation of the data needed to execute the algorithm efficiently. All of the DSP instructions control the multipliers 504A-504B, adders 510A-510C, compressor 506 and the accumulator 512, the functional units of each signal processing unit 300A-300D.

[0059] In the 40 bit instruction word, the type of extension from the 20 bit instruction word falls into five categories:

[0060] 1) Control and Specifier extensions that override the control bits in mode registers

[0061] 2) Type extensions that override the type specifier in address registers

[0062] 3) Permute extensions that override the permute specifier for vector data in address registers

[0063] 4) Offset extensions that can replace or extend the offsets specified in the address registers

[0064] 5) DSP extensions that control the lower rows of functional units within a signal processing unit 300 to accelerate block processing.

[0065] The 40-bit control instructions with the 20 bit extensions further allow a large immediate value (16 to 20 bits) to be specified in the instruction and powerful bit manipulation instructions.

[0066] Efficient DSP execution is provided with 2x20-bit DSP instructions with the first 20-bits controlling the top functional units (adders 501A and 510B, multiplier 504A, compressor 506) that interface to data buses from memory and the second 20 bits controlling the bottom functional units (adder 510C and multiplier 504B) that use internal or local data as operands. The top functional units, also referred to as main units, reduce the inner loop cycles in the inner loop 602 by parallelizing across consecutive taps or sections. The bottom functional units cut the outer loop cycles in the outer loop 601 in half by parallelizing block DSP algorithms across consecutive samples.

[0067] Efficient DSP execution is also improved by the hardware architecture of the present invention. In this case, efficiency is improved in the manner that data is supplied to and from data memory 202 to feed the four signal processing units 300 and the DSP functional units therein. The data highway is comprised of two buses, X bus 531 and Y bus 533, for X and Y source operands, and one Z bus 532 for a result write. All buses, including X bus 531, Y bus 533, and Z bus 532, are preferably 64 bits wide. The buses are uni-directional to simplify the physical design and reduce transit times of data. In the preferred embodiment when in a 20 bit DSP mode, if the X and Y buses are both carrying operands read from memory for parallel execution in a signal processing unit 300, the parallel load field can only access registers within the register file 413 of the RISC control unit 302. Additionally, the four signal processing units 300A-300D in parallel provide four parallel MAC units (multiplier 504A, adder 510A, and accumulator 512) that can make simultaneous computations. This reduces the cycle count from 4 cycles ordinarily required to perform four MACs to only one cycle.

Dyadic DSP Instructions

[0068] All DSP instructions of the instruction set architecture of the ASSP 150 are dyadic DSP instructions within the 20 bit or 40 bit instruction word. A dyadic DSP instruction informs the ASSP in one instruction and one cycle to perform two operations. Referring now to FIG. 6B is a chart illustrating the permutations of the dyadic DSP instructions. The dyadic DSP instruction 610 includes a main DSP operation 611 (MAIN OP) and a sub DSP operation 612

(SUB OP), a combination of two DSP instructions or operations in one dyadic instruction. Generally, the instruction set architecture of the present invention can be generalized to combining any pair of basic DSP operations to provide very powerful dyadic instruction combinations. Compound DSP operational instructions can provide uniform acceleration for a wide variety of DSP algorithms not just multiply-accumulate intensive filters. The DSP instructions or operations in the preferred embodiment include a multiply instruction (MULT), an addition instruction (ADD), a minimize/maximize instruction (MIN/MAX) also referred to as an extrema instruction, and a no operation instruction (NOP) each having an associated operation code (“opcode”). Any two DSP instructions can be combined together to form a dyadic DSP instruction. The NOP instruction is used for the MAIN OP or SUB OP when a single DSP operation is desired to be executed by the dyadic DSP instruction. There are variations of the general DSP instructions such as vector and scalar operations of multiplication or addition, positive or negative multiplication, and positive or negative addition (i.e. subtraction)

[0069] Referring now to FIG. 6C and FIG. 6D, bitmap syntax for an exemplary dyadic DSP instruction is illustrated. FIG. 6C illustrates bitmap syntax for a control extended dyadic DSP instruction while FIG. 6D illustrates bitmap syntax for a non-extended dyadic DSP instruction. In the non-extended bitmap syntax the instruction word is the twenty most significant bits of a forty bit word while the extended bitmap syntax has an instruction word of forty bits. The three most significant bits (MSBs), bits numbered 37 through 39, in each indicate the MAIN OP instruction type while the SUB OP is located near the middle or end of the instruction bits at bits numbered 20 through 22. In the preferred embodiment, the MAIN OP instruction codes are 000 for NOP, 101 for ADD, 110 for MIN/MAX, and 100 for MULT. The SUB OP code for the given DSP instruction varies according to what MAIN OP code is selected. In the case of MULT as the MAIN OP, the SUB OPs are 000 for NOP, 001 or 010 for ADD, 100 or 011 for a negative ADD or subtraction, 101 or 110 for MIN, and 111 for MAX. In the preferred embodiment, the MAIN OP and the SUB OP are not the same DSP instruction although alterations to the hardware functional blocks could accommodate it. The lower twenty bits of the control extended dyadic DSP instruction, the extended bits, control the signal processing unit to perform rounding, limiting, absolute value of inputs for SUB OP, or a global MIN/MAX operation with a register value.

[0070] The bitmap syntax of the dyadic DSP instruction can be converted into text syntax for program coding. Using the multiplication or MULT non-extended instruction as an example, its text syntax for multiplication or MULT is

(vmul|vmuln).(vadd|vsub|vmax|sadd|ssub|smax) da,
sx, sa, syl,(ps0)|ps1]

The “vmul|vmuln” field refers to either positive vector multiplication or negative vector multiplication being selected as the MAIN OP. The next field, “vadd|vsub|vmax|sadd|ssub|smax”, refers to either vector add, vector subtract, vector maximum, scalar add, scalar subtraction, or scalar maximum being selected as the SUB OP. The next field, “da”, refers to selecting one of the registers within the accumulator for storage of results. The field “sx” refers to selecting a register within the RISC

register file **413** which points to a memory location in memory as one of the sources of operands. The field “sa” refers to selecting the contents of a register within the accumulator as one of the sources of operands. The field “sy” refers to selecting a register within the RISC register file **413** which points to a memory location in memory as another one of the sources of operands. The field of “[,(ps0)|ps1]” refers to pair selection of keyword PS0 or PS1 specifying which are the source-destination pairs of a parallel-store control register. Referring now to **FIG. 6E and 6F**, lists of the set of 20-bit DSP and control instructions for the ISA of the present invention is illustrated. **FIG. 6G** lists the set of extended control instructions for the ISA of the present invention. **FIG. 6H** lists the set of 40-bit DSP instructions for the ISA of the present invention. **FIG. 6I** lists the set of addressing instructions for the ISA of the present invention.

[0071] Referring now to **FIG. 7**, a block diagram illustrates the instruction decoding for configuring the blocks of the signal processing unit **300**. The signal processor **300** includes the final decoders **704A** through **704N**, and multiplexers **720A** through **720N**. The multiplexers **720A** through **720N** are representative of the multiplexers **514**, **516**, **520**, and **522** in **FIG. 5B**. The predecoding **702** is provided by the RISC control unit **302** and the pipe control **304**. An instruction is provided to the predecoding **702** such as a dyadic DSP instruction **600**. The predecoding **702** provides preliminary signals to the appropriate final decoders **704A** through **704N** on how the multiplexers **720A** through **720N** are to be selected for the given instruction. Referring back to **FIG. 5B**, in a dyadic DSP instruction the MAIN OP generally, if not a NOP, is performed by the blocks of the multiplier **M1504A**, compressor **506**, adder **A1510A**, and adder **A2510B**. The result is stored in one of the registers within the accumulator register **AR 512**. In the dyadic DSP instruction the SUB OP generally, if not a NOP, is performed by the blocks of the adder **A3510C** and the multiplier **M2504B**. For example, if the dyadic DSP instruction is to perform is an ADD and MULT, then the ADD operation of the MAIN OP is performed by the adder **A1510A** and the SUB OP is performed by the multiplier **M1504A**. The predecoding **720** and the final decoders **704A** through **704N** appropriately select the respective multiplexers **720A** through **720B** to select the MAIN OP to be performed by the adder **A1510A** and the SUB OP to be performed by the multiplier **M2504B**. In the exemplary case, multiplexer **520A** selects inputs from the data typer and aligner **502** in order for adder **A1510A** to perform the ADD operation, multiplexer **522** selects the output from adder **510A** for accumulation in the accumulator **512**, and multiplexer **514B** selects outputs from the accumulator **512** as its inputs to perform the MULT SUB OP. The MAIN OP and SUB OP can be either executed sequentially (i.e. serial execution on parallel words) or in parallel (i.e. parallel execution on parallel words). If implemented sequentially, the result of the MAIN OP may be an operand of the SUB OP. The final decoders **704A** through **704N** have their own control logic to properly time the sequence of multiplexer selection for each element of the signal processor **300** to match the pipeline execution of how the MAIN OP and SUB OP are executed, including sequential or parallel execution. The RISC control unit **302** and the pipe control **304** in conjunction with the final decoders **704A** through **704N** pipelines instruction execution by pipelining the instruction itself and by providing pipelined control

signals. This allows for the data path to be reconfigured by the software instructions each cycle.

[0072] As those of ordinary skill will recognize, the present invention has many advantages. One advantage of the present invention is that the ISA is adapted to DSP algorithmic structures providing compact hardware to consume low-power which can be scaled to higher computational requirements. Another advantage of the present invention is that the signal processing units have direct access to operands in memory to reduce processing overhead associated with load and store instructions.

[0073] Another advantage of the present invention is that pipelined instruction execution is provided so that instructions may be issued every cycle. Another advantage of the present invention is that the signal processing units can be configured cycle by cycle.

[0074] The preferred embodiments of the present invention are thus described. While the present invention has been described in particular embodiments, it may be implemented in hardware, software, firmware or a combination thereof and utilized in systems, subsystems, components or sub-components thereof. When implemented in software, the elements of the present invention are essentially the code segments to perform the necessary tasks. The program or code segments can be stored in a processor readable medium or transmitted by a computer data signal embodied in a carrier wave over a transmission medium or communication link. The “processor readable medium” may include any medium that can store or transfer information. Examples of the processor readable medium include an electronic circuit, a semiconductor memory device, a ROM, a flash memory, an erasable ROM (EROM), a floppy diskette, a CD-ROM, an optical disk, a hard disk, a fiber optic medium, a radio frequency (RF) link, etc. The computer data signal may include any signal that can propagate over a transmission medium such as electronic network channels, optical fibers, air, electromagnetic, RF links, etc. The code segments may be downloaded via computer networks such as the Internet, Intranet, etc. In any case, the present invention should not be construed as limited by such embodiments, but rather construed according to the claims that follow below.

1-20. (canceled)

21. A method of instruction execution for a signal processor, the method comprising:

writing a dyadic DSP instruction having a main operation and a sub operation into a memory;

reading the dyadic DSP instruction from the memory to begin execution of the main operation and the sub operation; and

addressing at least three operands for execution with the main operation and the sub operation.

22. The method of claim 21, further comprising:

executing the main operation and the sub operation using the at least three operands.

23. The method of claim 22, further comprising:

storing the results of the execution of the dyadic DSP instruction.

24. The method of claim 21, wherein,
the addressing addresses a register for one source of one of the at least three operands.

25. The method of claim 21, wherein,
the addressing addresses a register that contains a pointer to an address into the memory for one source of one of the at least three operands.

26. The method of claim 21, wherein,
the addressing addresses the memory for one source of one of the at least three operands.

27. The method of claim 21, wherein,
the main operation and the sub operation are two operations selected from the set of multiplication, addition, comparison with a minimum or maximum value, and no operation.

28. A method in a signal processor, the method comprising: in response to a dyadic DSP instruction,
selecting input and output data paths between a plurality of DSP functional blocks;
executing a main operation of the dyadic DSP instruction on a first pair of operands using the plurality of DSP functional blocks;
executing a sub operation of the dyadic DSP instruction on a second pair of operands using the plurality of DSP functional blocks; and

storing the results of the execution of the main operation and the sub operation of the dyadic DSP instruction.

29. The method of claim 28, further comprising:

prior to the executing, addressing at least three operands for execution with the main operation and the sub operation.

30. The method of claim 29, wherein,

one of the at least three operands for execution is a result of the execution of the main operation or the sub operation.

31. The method of claim 28, wherein,

the execution of the main operation and the sub operation are performed in parallel during the same cycle.

32. The method of claim 28, wherein,

the execution of the main operation and the sub operation are performed sequentially during different cycles.

33. The method of claim 28, wherein,

the main operation and the sub operation are two operations selected from the set of multiplication, addition, comparison with a minimum or maximum value, and no operation.

* * * * *