Description

CROSS-REFERENCE TO RELATED APPLICATION

This application contains subject matter related to the invention disclosed and claimed in the European Patent Application 87110756.1 entitled ACCELERATED VALIDITY RESPONSE PERMITTING EARLY ISSUE OF INSTRUCTIONS DEPENDING UPON OUTCOME OF FLOATING POINT OPERATIONS.

BACKGROUND OF THE INVENTION

This invention is in the field of machine emulation and especially concerns a source which emulates operation of a source CPU by translating source CPU instructions into target CPU instructions for issue and execution by a target CPU. More particularly, the invention relates to extensions of the emulation ability of such a system through the provision of means for accelerating the translating and issuing functions when source floating point arithmetic instructions are being emulated and which emulate the execution of source floating point RX-type instructions by execution of target register-to-register floating point arithmetic instructions.

Emulation is the imitation of the operation of a first ("source") CPU by a second ("target") CPU. The target CPU is specially programmed and architected to permit it to execute programs written for the source CPU. A program written for the source CPU comprises a sequence of source instructions which are provided, one-by-one, to the target CPU. The target CPU responds to each source instruction by executing one or more target instructions.

In U. S. Patent No. 4,587,612 of Fisk et al., assigned to the present assignee, an emulation assist processor (EAP) receives a source instruction stream and maps each source instruction to one or more target instructions, the target instructions being passed to an instruction processing unit (IPU) of the target CPU. In the incorporated patent, the EAP converts multi-field source instructions into multi-field target instructions and streams the target instructions to the IPU for processing and issuing.

As is known, when the source CPU comprises a machine such as an IBM 370 host computer (described in U. S. Patent No. 3,400,371, assigned to the present assignee,), the source instruction set includes multi-field floating point arithmetic instructions, primarily of the RX-type. Source instruction programs utilizing the IBM 370 instruction set characteristically are conditioned by the state of a condition code (CC), an indicator which is set according to the outcome of certain instruction among which are floating point arithmetic instructions. When an IBM 370 floating point instruction produces an abnormal outcome (such as an all zero result) or attempts an abnormal operation (divide by zero), an interrupt indicator is set which transfers control from the executing program to a supervisory for certain interrupt procedures. Emulation of the branching and interrupt features of an IBM 370 source program requires that the target program maintains condition code and interrupt indicators to effectively map program branches and interrupts.

While conceding that the branching and interrupt correspondence between source and target programs must be maintained, it is recognized that the speed of emulation can be enhanced by the ability to reliably predict the state of the condition code and interrupt indicators before the completion of executing floating point arithmetic operations. However, any such enhancement must account not only for the translation of source to target instructions, but also must take into account the issuance of converted target instructions.

One of the potential operational environments of the EAP of U. S. Patent No. 4,587,612 imposes certain architectural bottlenecks to execution of target floating point instructions. In this regard, the instruction unit and floating point unit of the target CPU are interconnected by a 32-bit wide databus. When source RX-type floating point instructions are emulated, an operand denoted in the X field of the source instruction must be fetched from memory for provision to the FPU. Passage of the responsibility for obtaining the X-field operand from memory lengthens the instruction execution time. Execution of a target floating point instruction is further lengthened when the source instruction is an extended RX-type requiring two sequential memory accesses to obtain a 64-bit operand over the 32-bit databus.

THE INVENTION

The invention has application in a system for emulating the execution of source CPU instructions, which system includes a target CPU embracing a memory, a number of functional units, and a target instruction unit which processes and issues translated target instructions for execution by target functional units. The emulation system also includes an instruction translating unit (translator) for receiving and converting source instructions to target instructions. The invention resides in an improvement to the system and has the form of a mechanism for accelerating the translation and execution of floating point instructions. In this regard, the invention includes a token generator in the

translator which generates a floating point wait token signal in response to translation of a source floating point instruction. An outcome generator in a floating point functional unit of the target CPU produces completion signals indicative of the outcome of issued target floating point instructions in advance of completion of the execution of the issued instructions. The system includes system means for restraining the translation and issue of instructions in response to the token signal. Finally, a multi-stage wait queue located in the translator and connected to the system means receives and retains the wait token signal when it is generated and extinguishes the token signal in response to the completion signals indicating the outcome of the target floating point instruction translated from the source floating point instruction.

The invention stems from a number of unanticipated observations. The inventors have observed that: if instruction translation and target instruction issue are responsive to an accelerated indication of the execution results of target floating point instructions, translation and issuing, and therefore emulation, can be accelerated; if a source RX-type instruction X-field operand can be obtained prior to translated target instruction issue, a source RX floating point instruction can be emulated by execution of a target RR instruction, thereby obviating a memory access cycle after instruction issue; and if the two halves of the X-field operand of an extended floating point source instruction are obtained on successive memory access cycles, emulation of the source instruction can be emulated as a set of target floating point instructions which transfer the operand and initiate execution of an RR target instruction sequence which reduces the number of cycles required to emulate the source instruction.

Therefore, the primary objective of the present invention is to improve the operation of a system which emulates the performance of a source CPU program that includes floating point arithmetic instructions of the RX-type by accelerating the translation of source to target instructions and the issue of converted target instructions.

A further objective of the present invention is to improve the operation of such an emulation system by reducing the time required for target CPU operations that emulate execution of source RX-type floating point arithmetic instructions.

These and other objects and further attendant advantages of the subject invention will become more evident when the detailed description is read in connection with the below-described drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 illustrates the architectural arrangement of a system for emulating a program provided by a source CPU of the IBM 370 type.

Fig. 2 illustrates in greater detail functional elements of the emulator-assist processor (EAP) and instruction processing unit (IPU) which accelerate the translation and target issuing functions in response to accelerated completion indications provided by a floating point unit (FPU) prior to the completion of target floating point arithmetic operations.

Fig. 3 illustrates the functional elements included in the IPU and FPU that support the emulation of RX-type floating point source instructions by RR-type floating point target instructions.

Fig. 4 is a flow diagram illustrating the sequence of operations executed by the arrangement of Fig. 3 in emulating RX short source instructions.

Fig. 5 is a flow diagram illustrating the sequence of operations executed by the arrangement of Fig. 3 in emulating extended RX-type floating point source instructions.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Fig.1 illustrates the application environment of the invention. This environment includes a reduced instruction set computer (RISC) which forms the target CPU. As described by Radin in "The 801 Minicomputer," published in the IBM Journal of Research and Development, Volume 27, No. 3, May 1983, at pages 237-246, such a computer includes separate instruction and data streams sourced by respective caches. In this computer, instructions are obtained from an instruction cache and data from a separate data cache, both of which are located in and managed by an instruction processing unit (IPU) 10. The contents of the caches are replenished by IPU access of a memory 12. The IPU 10 processes cached instructions and issues them for execution by one or more processing units (PU) 14 and a floating point processing unit (FPU) 16. Since the units 10, 12, 14, and 16 form the target CPU, the instructions issued by the IPU 10 are termed target instructions and are issued to the processing units on a target instruction bus 19. Operand and result (O/R) data is exchanged between the IPU 10 and the processing units 14 and 16 on the O/R databus 18.

In the postulated environment, the exemplary 801 computer has a 32-bit architecture which is reflected in a corresponding width for the data and instruction buses 18 and 19. Further, in addition to the data cache, the IPU 10 of the 801 includes a bank of general-purpose registers.

When used as an emulator, the target CPU operates in conjunction with an emulator-assist processor (EAP) 20. In the emulation mode, a program consisting of instructions from a source CPU is streamed through a portion of the data cache in the IPU 10 on the data path 21, 10, 22 to the EAP 20. In the EAP 20, the source instructions are converted into target instructions and inserted into a target CPU instruction stream which does not disturb the normal target CPU instruction execution sequence executed by the IPU 10. The converted instruction stream is provided to the IPU 10 on datapath 23. The cooperative operations of the IPU 10 and EAP 20 in performing instruction conversion and generating the target instruction stream are well explained in the incorporated Fisk Patent and will not be repeated here. However, certain of the structures and functions involved in the instruction conversion must be set out in order to adequately explain the invention.

Refer now to Fig.2, where the EAP 20, IPU 10, and FPU 16 are shown. The IPU is shown in two functional sections 10a and 10b only to facilitate the following discussion; the IPU 10 in fact comprises a single unit. In Fig.2, the portion 10a of the IPU 10 which receives the source instruction stream on 21 includes a data cache 24 which terminates the source stream, and an instruction cache 25 which contains microinstructions. The source instruction stream is provided to the EAP 20 from the data cache on signal line 22 and staged, instruction by instruction, through a source instruction register (SIR) 26 which provides a platform against which an instruction mapping circuit 27 can operate. As explained in the Fisk patent, the instruction mapping circuit 27 generates a next instruction address which is provided to the instruction cache 25 in the IPU 10, and which results in the provision of a microinstruction from the instruction cache, the microinstruction being provided to the EAP on the signal line 30. The microinstruction on the signal line 30 is staged to the EAP in the microinstruction register (MIR) 31. As is taught in the Fisk patent, the microinstruction held in the register 31 also includes a next instruction field which can generate the next instruction address on line 29. Thus, a source instruction in the register 26 can result in the generation of one or a sequence of microinstructions which are fed to the EAP on the signal line 30. Relatedly, emulation of a source instruction can involve retrieval of a microinstruction sequence from the instruction cache, and execution of a resulting target instruction sequence by the target processor. The microinstruction provided to the register 31 consists of a control section and a skeleton target instruction. The skeleton target instruction is provided on signal line 23a to an instruction merge register 32 in the IPU 10. As taught in the Fisk patent, the skeleton instruction consists of an OP code field and register and displacement fields. The control portion of the microinstruction in the register 31 and the operand fields of the source instruction in the register 26 are used by the mapping circuit 27 to fill in zeroed operand and/or control fields of the skeleton instruction in an instruction merge register 32 in the IPU 10. The fill information is provided from the mapping circuit on the signal line 23b. Thus, the signal lines 23a and 23b form the signal line 23 of Fig.1 on which the target instruction stream is provided to the IPU 10. When an instruction is issued by the IPU 10, it is transferred from the instruction merge register 32 to an instruction register (I REG) 33. When the target instruction resides in the I register 33, the IPU 10 undertakes a course of instruction issue operations, summarized below.

The operational context of the invention assumes a source instruction stream including multifield instructions such as are included in the IBM 370 CPU instruction set. At least one distinctive format in this set is termed the RX (register/index) format. The RX format has the following form:

OP R_1 X_2 B_2 D_2

In an RX instruction, the operation is indicated by the OP code field. A first operand is located in register R_1, while the second operand is at a main memory location having the address $X_2 + B_2 + D_2$. X_2 and B_2 refer to general registers functioning, for the purpose of the instruction, as index and base registers, while D_2 is a displacement. When the instruction is executed, the result is placed in register R_1.

As is known, RX-type instructions include a subset of instructions involving arithmetic operations that are performed on floating point operands. Floating point arithmetic is well-understood. Floating point data format, terminology, and operations are reviewed at length in the article by Anderson et al. entitled "The IBM System/360 Model 91: Floating-Point Execution Unit," which appeared in the IBM Journal dated January 1967 at pages 34-53 and in the cross-referenced patent application.

RX-type floating point arithmetic source instructions are converted into floating point arithmetic target instructions by the EAP 20 for issue by the IPU 10 and execution by the FPU 16.

An RX-type floating point arithmetic instruction can produce a result which alters the condition code (CC) or which raises a program interrupt request (IR). As is known, the state of the CC conditions certain branches in the executing source program. An interrupt request will result in transfer of program control from the executing program. Therefore, these indicators must be accounted for in the emulation of a source instruction stream. In the EAP 20, the current state of the CC and IR are maintained in a status register 34. The instruction mapping circuit 27 inspects the contents of the status register 34 in order to undertake branching and interrupt activities, when necessary.

The CC and IR are provided by, among other sources, the FPU 16, where floating point instructions which affect the indicators are performed. The CC and IR are provided on the signal lines 35 and 36, respectively. Also provided by the FPU 16 is an accelerated validity response (AVR) signal on a signal line 38. The AVR is raised by the FPU 16 at the time when the CC and IR validly reflect the outcome of a floating point arithmetic instruction currently being executed by the FPU 16. As taught in the U. S. Patent Application, the FPU 16 can set the CC and IR in advance of the completion of a floating point arithmetic instruction whose result determines their states. When the CC and IR have been generated, the FPU 16 raises the AVR, signalling to the EAP 20 that the contents of the register 34 can be inspected. The time of occurrence of the AVR can range from soon after commencement of a floating point arithmetic instruction to the time of completion of an instruction, which is the time at which the CC and IR are conventionally sampled.

In order to take advantage of the accelerated production of the CC and IR by the FPU 16, the invention provides the EAP 20 with a wait token circuit which notifies the EAP 20 and IPU 10 when a floating point arithmetic target instruction has been translated and issued, restrains further translation and issuance of target instructions until production of the CC and IR in response to execution of the issued floating point arithmetic target instruction, and notifies the EAP 20 and IPU 10 when the CC and IR have been produced, permitting the early resumption of instruction translation and issue.

When a source floating point arithmetic instruction is contained in the register 26, the instruction mapping circuit 27 generates the first address of a microinstruction sequence adapted for translating the source instruction. All microinstruction sequences generated as the result of a floating point arithmetic instruction terminate with a WAIT microinstruction. The WAIT microinstruction is the signal to the invention to begin production of a wait token. In this regard, when the source floating point arithmetic instruction is registered at 26, its OP code field is provided to a conventional decoder (D) 40 which activates a first token precursor signal. When the WAIT microinstruction at the end of the stimulated sequence is registered at 31, the OP code of the microinstruction is also provided to the decoder 40 which generates a second token precursor signal. A coincidence detector (CD) 42 responds to the two token precursor signals by activating a WAIT TOKEN signal on signal line 43. The WAIT TOKEN signal is provided to a wait token queue 45, consisting of a multi-stage latch sequence including latches 46 and 47. When the WAIT microinstruction is in the microinstruction register 31, a skeleton target floating point instruction is entered under the instruction merge register 32 on the signal line 23a, with the operand field data for the instruction being provided by the instruction mapping circuit 27 on signal line 23b. The WAIT TOKEN signal is clocked from the latch 46 to the latch 47 by one phase (CLK A) of a multi-phase pipeline clock. When the WAIT TOKEN signal enters the latch 47, the latch activates a HOLD signal that is provided to the instruction mapping circuit 27, the instruction merge and instruction registers 32 and 33, and to the status register 34. The HOLD signal prevents alteration of data in the registers 32, 33, and 34 for so long as the HOLD signal is active. In addition, the HOLD signal suspends the operation of the instruction mapping circuit 27.

Concurrent with the transfer of the WAIT TOKEN signal into the latch 47 in response to a second multi-phase clock (CLK B), the target floating point instruction shifts from the instruction merge to the instruction register 33 in the IPU 10. In the instruction register 33, the instruction issues and is provided to the FPU 16. The FPU 16 undertakes operations for executing the issued instruction and also operates to set the CC and IR to the states reflecting the outcome of the executed instruction. When the CC and IR have been produced by the FPU 16, the AVR is raised, which clears the latch 47, thereby deactivating the HOLD signal.

When the HOLD signal is deactivated, the current CC and IR are entered into the status register 34 from the FPU 16. When the entry is made, the instruction mapping circuit 27 inspects the contents of the status register 34 to undertake any required branching or interrupt operations. Following deactivation of the HOLD signal and the generation of the next skeleton instruction, a merged instruction will lodge in the register 32 and be transferred for issue to the instruction register 33.

It will be evident to those skilled in the art that the WAIT TOKEN circuit 45 responds to acceleration of the AVR and permits target instruction con-

version and issue to take place in advance of the completion of an executing target floating point arithmetic instruction. The skilled artisan will appreciate that this accelerated operation can only enhance the concurrency of operations in the emulator of Fig. 1, thereby increasing the overall speed of its operations.

Reference to Fig.3 will engender an understanding of how the invention achieves the second of the stated objectives: that of reducing the time required for execution of target machine operations that emulate execution of source RX-type floating point arithmetic instructions. In Fig.3, the IPU 10 includes, in addition to the instruction register 33, a bank of general purpose registers (GPR) 50, a decode/control circuit 52, and a pair of multiplexers (MUX) 54, 56. The general purpose registers are used for instruction processing operations in the IPU 10. The decode/control circuit conventionally receives OP code information from the OP field of the target instruction in the register 33 and information in the D field of the instruction. The D field specifies a control operation to be executed and identifies the controller responsible for the execution. When designated, the decode/control circuit 52 generates control signals for the MUXES 54, 56 and the GPR's 50 to execute the OP code indicated instruction. When an instruction processing operation is required, the D field identifies the decode/control circuit 52 and enables the circuit to conduct a specified operation. Such an operation includes data fetched from memory, entailing operation of the MUX 54 to conduct the fetched data to a designated GPR. When data is to be exchanged with the O/R databus 18, the circuit 52 operates the MUX 56 to conduct the data and also operates a GPR to source or sink the data. The complement of operations performed by the decode/control circuit 52 also includes a LOAD IMMEDIATE instruction which will cause a value in one of the register fields (RT, RA, or RB) of the target instruction in the register 33 to be entered into a predetermined one of the GPR registers.

Fig.3 also illustrates, in greater detail, the FPU 16. The FPU 16 includes a control unit 60 which can be designated by a D field of an issued target instruction and which can control and synchronize the operation of FPU resources to conduct the execution of an instruction specified in the OP and D fields of the issued target instruction. The FPU has internal resources including a bank of floating point registers (FPR) 62. The FPR's 62 comprise a plurality of general purpose 64-bit registers for temporary storage of the operands and results of floating point operations conducted by the FPU 16. FPU arithmetic and logic operations are conducted by execution units (X) 64 and 65. As is known, such units can include, for example, a multiply/divide execution unit, an add/subtract/shift unit, and a unit for performing radical operations. Since the IPU and its other associated functional units are assembled in a 32-bit data and instruction architecture, the data interface of the FPU 16 is a 32-bit interface register (IR) 66, which exchanges data with the 32-bit databus 18.

Referring now to Figs.3 and 4 together, the operation of the invention in emulating "short" RX-type floating point arithmetic source instructions can be understood. The term "short" indicates that the operands of an RX-type instruction are 32 bits in length. Thus, the data interface 66 between the IPU 10 and FPU 16 can transfer an entire operand result in a single cycle. In step S1 of Fig. 4, when a short source RX-type instruction enters the EAP 20, the EAP will undertake a microinstruction sequence which includes provision of a LOAD IMMEDIATE instruction to the IPU 10, designating the IPU as the unit of execution. The instruction will carry in one of its register fields (RT, RA, or RB) the memory address corresponding to the sum of the X_2, B_2, and D_2 fields of the RX-type instruction. The address will be loaded into GPR_i. Next, during the microinstruction sequence, a LOAD instruction to the IPU will result in the operand stored at the address held in GPR_i being loaded into GPR_j. This is step S2 in Fig.4. Now, the final instruction of the microinstruction sequence results in a target floating point arithmetic instruction being loaded into the register 33 for issue. The issued instruction is referred to as an extended-RX (ERX) short-type target instruction. Such an instruction results in the IPU control circuit 52 placing the contents of GPR_j on the databus 18; this is the X-field operand of the source RX instruction. The FPU receives and decodes the issued instruction as a register-to-register (RR) floating point arithmetic instruction in which the first operand is located in the FPR designated by the RA field of the instruction and the second operand is located in the interface register 66, which is designated in the RB field of the issued instruction. The result of the target instruction execution is placed in the FPR denoted in the RT field of the instruction.

The procedure of Fig.4 is an efficient one for emulating source RX-type floating point arithmetic instructions by execution of target register-to-register floating point instructions. Without the RX enhancement of Fig.4, it would be necessary to insert the added step of transferring the operand held in GPR_j to a floating point register and then executing a floating point register-to-register floating point instruction. The extra step would be inserted between steps S2 and S3 of Fig.4.

Reference now to Figs.3 and 5 will provide an understanding of how the invention operates to emulate a source RX floating point arithmetic in-

struction designating a long (64-bit) operand. In the microinstruction sequence undertaken by the EAP in response to an extended RX-type instruction, the sequence of Fig.4 is initially followed by providing, in step S10, a LOAD IMMEDIATE target instruction for execution by the IPU, which places the X field address of the RX instruction into GPR_i. It will be appreciated that the address in GPR_i defines a double-word operand comprising 64 bits. Thus, following the LOAD IMMEDIATE instruction, in steps S11 and S12, two LOAD instructions are sent in sequence, causing the IPU to first load the high order (HO) four bytes of the X-field operand into GPR_j and then the low order (LO) four bytes of the operand into GPR_k. In step S13, the penultimate instruction provided to the IPU as a result of the microinstruction sequence undertaken by the EAP consists of an extended RX (ERX) long-type floating point instruction to be issued to the FPU. Essentially, the instruction is recognized by the FPU as an extended register-to-register instruction involving a pair of 64-bit operands, one stored in the FPR designated by the RA field of the instruction. The FPU obtains the second operand from the interface register 66. The second operand is staged to the FPU through the interface register 66 by transferring first the high order four bytes in GPR_j concurrently with issue of the instruction in step S13, and then, in step S14, transferring the low order four bytes in response to the final target instruction resulting from the microinstruction sequence. The final instruction is denoted as FINRX and is executed only by the IPU. Instruction execution consists of transfer of the low order four bytes from GPR_k to the interface register 66. Thus, the execution sequence of Fig.5 permits efficient emulation of an extended source RX floating point instruction by means of an extended target register-to-register floating point operation. The FINRX dummy instruction adapts the emulation sequence to the 32-bit architecture of the target machine.

Claims

1. In a system for emulating the execution of source CPU instructions which includes a memory (12), a target instruction unit (10) which processes and issues translated target instructions, target functional units (14, 16) for executing target instructions, and an instruction translating unit (translator) (20) for receiving and converting source instructions to target instructions, an improvement for accelerating translation and execution of source instructions, said improvement being characterized by:
   means (40, 42) in said translator (20) for generating a floating point token signal in response to translation of a source floating point arithmetic instruction;
   outcome means in said floating point unit for generating completion signals (AVR) indicative of the outcome of the target floating point instruction issued in response to said translation, said completion signals being generated in advance of completing the execution of said issued target floating point instruction;
   a multi-stage wait queue means (45) in said translator connected to said system means for receiving and moving said token signal to a wait stage in response to said translation and for removing said wait token signal from said wait stage in response to said completion signals; and
   system means (10b, 27) connected to said wait stage for restraining the translation and issue of instructions for so long as said wait token signal is held in said wait stage.

2. The system of claim 1 wherein said source floating point arithmetic instruction is an RX-type instruction and said target instruction is an RR-type instruction executable against the operands of said RX instruction and further including:a register (50) in said instruction unit (10) for receiving a floating point operand indicated by the X-field of said RX instruction;
   a data interface means (66) in said floating point unit selectively connectable to said register for buffering said operand to said floating point unit; and
   means (52, 54, 56) in said instruction unit for obtaining said operand from a location in said memory indicated by said X-field prior to the issue of said target RR instruction and for connecting said data interface means to said register to transfer said operand concurrently with the issued target RR instruction.

3. The system of claim 1 wherein said source instruction is an extended RX-type instruction and said target instruction is an RR-type instruction, executable against the operands of said RX instruction and further including:
   first and second registers (50) in said instruction unit (10) for receiving respective first and second halves of an extended floating point operand indicated by the X-field of said RX instruction;
   a data interface means (66) in said floating point unit selectively connectable to said first or second register for sequentially buffering said first and second operand to said floating point unit; and
   means (52, 54, 56) in said instruction unit for obtaining the first and second halves of said

struction designating a long (64-bit) operand. In the microinstruction sequence undertaken by the EAP in response to an extended RX-type instruction, the sequence of Fig.4 is initially followed by providing, in step S10, a LOAD IMMEDIATE target instruction for execution by the IPU, which places the X field address of the RX instruction into GPR_i. It will be appreciated that the address in GPR_i defines a double-word operand comprising 64 bits. Thus, following the LOAD IMMEDIATE instruction, in steps S11 and S12, two LOAD instructions are sent in sequence, causing the IPU to first load the high order (HO) four bytes of the X-field operand into GPR_j and then the low order (LO) four bytes of the operand into GPR_k. In step S13, the penultimate instruction provided to the IPU as a result of the microinstruction sequence undertaken by the EAP consists of an extended RX (ERX) long-type floating point instruction to be issued to the FPU. Essentially, the instruction is recognized by the FPU as an extended register-to-register instruction involving a pair of 64-bit operands, one stored in the FPR designated by the RA field of the instruction. The FPU obtains the second operand from the interface register 66. The second operand is staged to the FPU through the interface register 66 by transferring first the high order four bytes in GPR_j concurrently with issue of the instruction in step S13, and then, in step S14, transferring the low order four bytes in response to the final target instruction resulting from the microinstruction sequence. The final instruction is denoted as FINRX and is executed only by the IPU. Instruction execution consists of transfer of the low order four bytes from GPR_k to the interface register 66. Thus, the execution sequence of Fig.5 permits efficient emulation of an extended source RX floating point instruction by means of an extended target register-to-register floating point operation. The FINRX dummy instruction adapts the emulation sequence to the 32-bit architecture of the target machine.

Claims

1. In a system for emulating the execution of source CPU instructions which includes a memory (12), a target instruction unit (10) which processes and issues translated target instructions, target functional units (14, 16) for executing target instructions, and an instruction translating unit (translator) (20) for receiving and converting source instructions to target instructions, an improvement for accelerating translation and execution of source instructions, said improvement being characterized by:
means (40, 42) in said translator (20) for generating a floating point token signal in response to translation of a source floating point arithmetic instruction;
outcome means in said floating point unit for generating completion signals (AVR) indicative of the outcome of the target floating point instruction issued in response to said translation, said completion signals being generated in advance of completing the execution of said issued target floating point instruction;
a multi-stage wait queue means (45) in said translator connected to said system means for receiving and moving said token signal to a wait stage in response to said translation and for removing said wait token signal from said wait stage in response to said completion signals; and
system means (10b, 27) connected to said wait stage for restraining the translation and issue of instructions for so long as said wait token signal is held in said wait stage.

2. The system of claim 1 wherein said source floating point arithmetic instruction is an RX-type instruction and said target instruction is an RR-type instruction executable against the operands of said RX instruction and further including:a register (50) in said instruction unit (10) for receiving a floating point operand indicated by the X-field of said RX instruction;
a data interface means (66) in said floating point unit selectively connectable to said register for buffering said operand to said floating point unit; and
means (52, 54, 56) in said instruction unit for obtaining said operand from a location in said memory indicated by said X-field prior to the issue of said target RR instruction and for connecting said data interface means to said register to transfer said operand concurrently with the issued target RR instruction.

3. The system of claim 1 wherein said source instruction is an extended RX-type instruction and said target instruction is an RR-type instruction, executable against the operands of said RX instruction and further including:
first and second registers (50) in said instruction unit (10) for receiving respective first and second halves of an extended floating point operand indicated by the X-field of said RX instruction;
a data interface means (66) in said floating point unit selectively connectable to said first or second register for sequentially buffering said first and second operand to said floating point unit; and
means (52, 54, 56) in said instruction unit for obtaining the first and second halves of said

operand from a location in said memory indicated by said X-field prior to the issue of said target RR instruction and for connecting said data interface means successively to said first and second registers to transfer said operand halves concurrently with the issue of said target RR instruction

4. A method for emulating execution of a source floating point RX-type instruction (source instruction) in an emulation system including a translator for translating said source instruction into an RR-type target floating point instruction (target instruction) issued by a target system instruction unit and executed by a target system floating point unit, comprising the steps of:
loading the operand indicated by the X field of said source instruction into a register located in said instruction unit;
issuing said target instruction to the floating point unit for execution against the operands of said source instruction;
concurrently with the issue of said target instruction,
- transferring said operand from said instruction unit register to said floating point unit,
- generating a floating point token signal in response to the translation of said source floating point arithmetic instruction;

generating completion signals indicative of the outcome of the target floating point instruction issued in response to said translation, said completion signals being generated in advance of completing the execution of said issued target floating point instruction;
receiving and moving said token signal to a wait stage in response to said translation and for removing said wait token signal from said wait stage in response to said completion signals; and
restraining the translation and issue of instructions for so long as said wait token signal is held in said wait stage.

5. A method for emulating the execution of an extended source floating point RX-type instruction (source instruction) in an emulation system including a translator for translating said source instruction into an RR-type target floating point instruction (target instruction) issued by a target system instruction unit and executed by a target system floating point unit, comprising the steps of:
loading one half of the operand indicated by the X-field of said source instruction into a first register located in the instruction unit;
loading the remaining half of said operand into a second register located in the instruction unit;
issuing said target instruction to the floating point unit for execution against the operands of said source instruction;
concurrently with the issue of said target instruction,
- transferring the operand half in said first register from said first register to said floating point unit;
- generating a floating point token signal in response to the translation of said source floating point arithmetic instruction;

transferring the operand half in said second register from said second register to said floating point unit;
generating completion signals indicative of the outcome of the target floating point instruction issued in response to said translation, said completion signals being generated in advance of completing the execution of said issued target floating point instruction;
receiving and moving said token signal to a wait stage in response to said translation and for removing said wait token signal from said wait stage in response to said completion signals; and
restraining the translation and issue of instructions for so long as said wait token signal is held in said wait stage.

Patentansprüche

1. Verbesserung zur Beschleunigung der Übersetzung und Ausführung von Quellbefehlen in einem System zur Emulation der Ausführung von CPU-Quellbefehlen, das einen Speicher (12), eine Zielbefehlseinheit (10) zur Verarbeitung und Abgabe von übersetzten Zielbefehlen, Zielfunktionseinheiten (14, 16) zur Ausführung von Zielbefehlen und eine Befehlsübersetzungseinheit (Übersetzer) (20) zum Empfang und zur Umsetzung von Quellbefehlen in Zielbefehle enthält, wobei diese Verbesserung gekennzeichnet ist durch:
Mittel (40, 42) im Übersetzer (20) zur Erzeugung eines Gleitkomma-Tokensignals als Antwort auf die Übersetzung eines arithmetischen Gleitkomma-Quellbefehls;
Ergebnismittel in der Gleitkommaeinheit zur Erzeugung von Beendigungssignalen (AVR), die das Ergebnis des als Antwort auf die Übersetzung abgegebenen Gleitkomma-Zielbefehls anzeigen und vor der Beendigung der Ausführung des abgegebenen Gleitkomma-Zielbefehls erzeugt werden;
ein mit den Systemelementen verbundenes Mittel für eine mehrstufige Warteschlange (45)

in diesem Übersetzer zum Empfang und zur Weiterleitung des Tokensignals hin zu einem Wartesignalpuffer als Antwort auf die Übersetzung und zum Löschen des Wartetokensignals aus diesem Wartesignalpuffer als Antwort auf das Beendigungssignal; und
mit dem Wartesignalpuffer verbundene Systemmittel (10b, 27) zur Unterdrückung der Übersetzung und Abgabe von Befehlen, solange das Wartetokensignal in dem Wartesignalpuffer gehalten wird.

2. Das System gemäß Anspruch 1, wobei der arithmetische Gleitkomma-Quellbefehl ein Befehl vom RX-Typ und der Zielbefehl ein auf die Operanden des RX-Befehls anwendbarer RR-Befehl ist, das ferner umfaßt:
ein Register (50) in der Befehlseinheit (10) zum Empfang eines vom X-Feld des RX-Befehls bezeichneten Gleitkommaoperanden;
ein Datenschnittstellenmittel (66) in der Gleitkommaeinheit, das selektiv mit diesem Register verbunden werden kann, um den Operanden in die Gleitkommaeinheit zu puffern; und
Mittel (52, 54, 56) in der Befehlseinheit zum Erhalt des Operanden von einem vom X-Feld bezeichneten Ort im Speicher vor der Abgabe des RR-Zielbefehls und zur Verbindung des Datenschnittstellenmittels mit dem erwähnten Register, um den Operanden gleichzeitig mit dem abgegebenen RR-Befehl zu übertragen.

3. Das System gemäß Anspruch 1, wobei der Quellbefehl ein erweiterter Befehl vom RX-Typ und der Zielbefehl ein auf die Operanden des RX-Befehls anwendbarer Befehl vom RR-Typ ist, das weiter umfaßt:
erste und zweite Register (50) in der Befehlseinheit (10) zum Empfang der ersten bzw. zweiten Hälfte eines vom X-Feld des RX-Befehls bezeichneten erweiterten Gleitkommaoperanden;
ein Datenschnittstellenmittel (66) in der Gleitkommaeinheit, die selektiv mit dem ersten oder zweiten Register verbunden werden kann, um den ersten und zweiten Operanden nacheinander in die Gleitkommaeinheit zu puffern; und
Mittel (52, 54, 56) in dieser Befehlseinheit zum Erhalt der ersten und zweiten Hälfte des Operanden von einem vom X-Feld bezeichneten Ort im Speicher vor der Abgabe des RR-Zielbefehls und zur aufeinanderfolgenden Verbindung des Datenschnittstellenmittels mit den ersten und zweiten Registern, um den Operanden gleichzeitig mit der Abgabe des RR-Befehls zu übertragen.

4. Verfahren zur Emulation der Ausführung eines Gleitkomma-Quellbefehls (Quellbefehl) vom RX-Typ in einem Emulationssystem mit einem darin enthaltenen Übersetzer zur Übersetzung des Quellbefehls in einen Gleitkomma-Zielbefehl (Zielbefehl) vom RR-Typ, der von einer Befehlseinheit eines Zielsystems abgegeben und von einer Gleitkommaeinheit eines Zielsystems ausgeführt wird, wobei das Verfahren folgende Schritte umfaßt:
Laden des vom X-Feld des Quellbefehls angezeigten Operanden in ein in der Befehlseinheit befindliches Register;
Abgabe des Zielbefehls an die Gleitkommaeinheit zur Ausführung mit den Operanden des Quellbefehls;
gleichzeitig mit der Abgabe des Zielbefehls
- Übertragung des Operanden von dem Register der Befehlseinheit an die Gleitkommaeinheit;
- Erzeugung eines Gleitkomma-Tokensignals als Antwort auf die Übersetzung des Gleitkomma-Arithmetik-Quellbefehls;

Erzeugung von Beendigungssignalen, die das Ergebnis des als Antwort auf die Übersetzung abgegebenen Gleitkomma-Zielbefehls anzeigen, wobei das Beendigungssignal vor der Beendigung der Ausführung des abgegebenen Gleitkomma-Zielbefehls erzeugt wird;
Empfang und Weiterleitung des Tokensignals hin zu einem Wartesignalpuffer als Antwort auf die Übersetzung und Löschen des Wartetokensignals aus diesem Wartesignalpuffer als Antwort auf das Beendigungssignal; und
Unterdrückung der Übersetzung und Abgabe von Befehlen, solange das Wartetokensignal in dem Wartesignalpuffer gehalten wird.

5. Verfahren zur Emulation der Ausführung eines erweiterten Gleitkomma-Quellbefehls (Quellbefehl) vom RX-Typ in einem Emulationssystem mit einem darin enthaltenen Übersetzer zur Übersetzung des Quellbefehls in einen Gleitkomma-Zielbefehl (Zielbefehl) vom RR-Typ, der von einer Befehlseinheit des Zielsystems abgegeben und von einer Gleitkommaeinheit des Zielsystems ausgeführt wird, wobei das Verfahren folgende Schritte umfaßt:
Laden einer Hälfte des vom X-Feld des Quellbefehls angezeigten Operanden in ein erstes Register, das sich in der Befehlseinheit befindet;
Laden der verbleibenden Hälfte dieses Operanden in ein zweites Register, das sich in der Befehlseinheit befindet;
Abgabe des Zielbefehls an die Gleitkommaeinheit zur Ausführung mit den Operanden des Quellbefehls;

gleichzeitig mit der Abgabe des Zielbefehls
- Übertragung der Operandenhälfte im ersten Register von diesem ersten Register an die Gleitkommaeinheit;
- Erzeugung eines Gleitkomma-Tokensignals als Antwort auf die Übersetzung des Gleitkomma-Arithmetik-Quellbefehls;

Übertragung der Operandenhälfte im zweiten Register von diesem zweiten Register an die Gleitkommaeinheit;
Erzeugung von Beendigungssignalen, die das Ergebnis des als Antwort auf die Übersetzung abgegebenen Gleitkomma-Zielbefehls anzeigen, wobei die Beendigungssignale vor der Beendigung der Ausführung des abgegebenen Gleitkomma-Zielbefehls erzeugt werden;
Empfang und Weiterleitung des Tokensignals hin zu einem Wartesignalpuffer als Antwort auf die Übersetzung und zum Löschen des Wartetokensignals aus dem Wartesignalpuffer als Antwort auf das Beendigungssignal; und
Unterdrückung der Übersetzung und Abgabe von Befehlen, solange das Wartetokensignal in dem Wartesignalpuffer gehalten wird.

Revendications

1. Dans un système pour émuler l'exécution d'instructions d'un CPU source qui comprend une mémoire (12), une unité d'instruction cible (10) qui traite et fournit des instructions cibles traduites, des unités fonctionnelles cibles (14, 16) pour exécuter des instructions cibles, et une unité de traduction d'instructions (traducteur) (20) pour recevoir et convertir des instructions sources en instructions cibles, un perfectionnement pour accélérer la traduction et l'exécution d'instructions sources, ledit perfectionnement étant caractérisé par:

des moyens (40, 42) dans ledit traducteur (20) pour générer un signal de jeton à virgule flottante en réponse à la traduction d'une instruction arithmétique source à virgule flottante;

un moyen de sortie dans ladite unité à virgule flottante pour générer des signaux d'achèvement (AVR) indicatifs de la sortie d'une instruction cible à virgule flottante en réponse à ladite traduction, lesdits signaux d'achèvement étant générés avant l'achèvement de l'exécution de ladite instruction cible à virgule flottante;

un moyen de file d'attente à plusieurs étages (45) dans ledit traducteur connecté auxdits moyens de système pour recevoir et transmettre ledit signal de jeton à un étage d'attente en réponse à ladite traduction et pour retirer ledit signal de jeton dudit étage d'attente en réponse auxdits signaux d'achèvement; et

des moyens de système (10b, 27) connectés audit étage d'attente pour arrêter la traduction et la délivrance des instructions aussi longtemps que ledit signal de jeton d'attente est maintenu dans ledit étage d'attente.

2. Système selon la revendication 1, dans lequel ladite instruction arithmétique source à virgule flottante est une instruction du type RX et ladite instruction cible est une instruction du type RR exécutable en fonction des opérandes de ladite instruction RX et comprenant en outre: un registre (50) dans ladite unité d'instruction (10) pour recevoir un opérande à virgule flottante indiqué par le champ X de ladite instruction RX;

un moyen d'interface de données (66) dans ladite unité à virgule flottante sélectivement connectable audit registre pour mémoriser temporairement ledit opérande vers ladite unité à virgule flottante; et

des moyens (52, 54, 56) dans ladite unité d'instruction pour obtenir ledit opérande à partir d'un emplacement de ladite mémoire indiqué par ledit champ X avant la sortie de ladite instruction cible RR et pour connecter ledit moyen d'interface de données audit registre de façon à transférer ledit opérande en même temps que l'instruction cible RR fournie.

3. Système selon la revendication 1, dans lequel ladite instruction source est une instruction du type RX étendue et ladite instruction cible est une instruction du type RR, exécutable en fonction des opérandes de ladite instruction RX et comprenant en outre:

un premier et un second registres (50) dans ladite unité d'instruction (10) pour recevoir les moitiés respectives d'un opérande à virgule flottante étendu indiqué par le champ X de ladite instruction RX;

un moyen d'interface de données (66) dans ladite unité à virgule flottante sélectivement connectable audit premier ou audit second registre pour séquentiellement mémoriser temporairement lesdits premier et second opérandes vers ladite unité à virgule flottante; et

des moyens (52, 54, 56) dans ladite unité d'instruction pour obtenir les première et seconde moitiés dudit opérande à partir d'un emplacement de ladite mémoire indiqué par ledit champ X avant la sortie de ladite instruction cible RR et pour connecter ledit moyen d'interface de données successivement auxdits premier et second registres de façon à transférer lesdites moitiés d'opérande en même temps que la sortie de ladite instruction cible RR.

4. Méthode pour émuler l'exécution d'une instruction source à virgule flottante du type RX (instruction source) dans un système d'émulation comprenant un traducteur pour traduire ladite instruction source en une instruction cible à virgule flottante du type RR (instruction cible) délivrée par une unité d'instruction cible du système et exécutée par une unité cible à virgule flottante du système, comprenant les étapes de:

charger l'opérande indiqué par le champ X de ladite instruction source dans un registre se trouvant dans ladite unité d'instruction;

délivrer ladite instruction cible à l'unité à virgule flottante pour exécution en fonction des opérandes de ladite instruction source;

transférer ledit opérande à partir dudit registre de l'unité d'instruction cible à ladite unité à virgule flottante en même temps que ladite instruction cible est délivrée;

générer un signal de jeton à virgule flottante en réponse à la traduction de ladite instruction arithmétique source à virgule flottante;

générer des signaux d'achèvement indicatifs de la sortie de l'instruction cible à virgule flottante fournie en réponse à ladite traduction, lesdits signaux d'achèvement étant générés avant l'achèvement de l'exécution de ladite instruction cible à virgule flottante;

transmettre ledit signal de jeton à un étage d'attente en réponse à ladite traduction et retirer ledit signal de jeton dudit étage d'attente en réponse auxdits signaux d'achèvement; et

arrêter la traduction et la délivrance des instructions aussi longtemps que ledit signal de jeton est maintenu dans ledit étage d'attente.

5. Méthode pour émuler l'exécution d'une instruction source à virgule flottante étendue du type RX (instruction source) dans un système d'émulation comprenant un traducteur pour traduire ladite instruction source en une instruction cible à virgule flottante du type RR (instruction cible) délivrée par une unité d'instruction cible de système et exécutée par une unité cible à virgule flottante de système, comprenant les étapes de:

charger une moitié de l'opérande indiqué par le champ X de ladite instruction source dans un premier registre se trouvant dans l'unité d'instruction;

charger la moitié restante dudit opérande dans un second registre se trouvant dans l'unité d'instruction;

délivrer ladite instruction cible à l'unité à virgule flottante pour exécution en fonction des opérandes de ladite instruction source;

transférer la moitié d'opérande dans ledit premier registre dudit premier registre à ladite unité à virgule flottante, en même temps que ladite instruction cible est délivrée.

générer un signal de jeton à virgule flottante en réponse à la traduction de ladite instruction arithmétique source à virgule flottante;

transférer la moitié d'opérande dans ledit second registre dudit second registre à ladite unité à virgule flottante;

générer des signaux d'achèvement indicatifs de la sortie de l'instruction cible à virgule flottante délivrée en réponse à ladite instruction, lesdits signaux d'achèvement étant générés avant l'achèvement de l'exécution de ladite instruction cible à virgule flottante;

recevoir et transmettre ledit signal de jeton à un étage d'attente en réponse à ladite traduction et retirer ledit signal de jeton dudit étage d'attente en réponse auxdits signaux d'achèvement; et

arrêter la traduction et la délivrance des instructions aussi longtemps que ledit signal de jeton est maintenu dans ledit étage d'attente.