



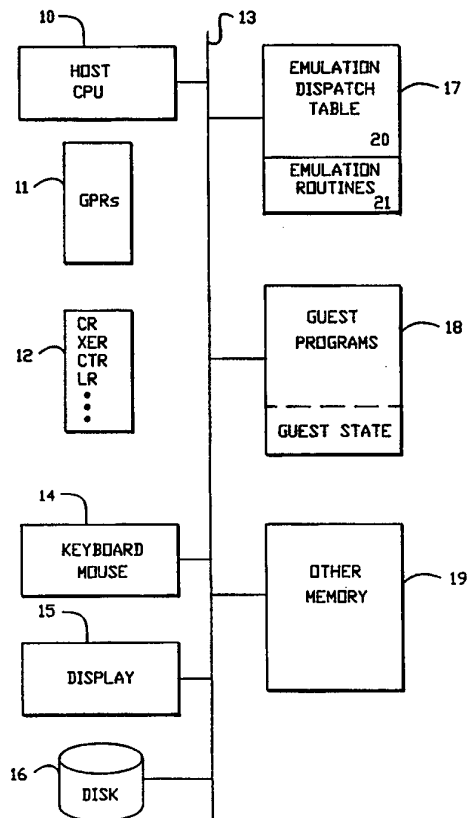
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification ⁵ : G06F 9/318, 9/455</p>	<p>A1</p>	<p>(11) International Publication Number: WO 94/27214 (43) International Publication Date: 24 November 1994 (24.11.94)</p>
<p>(21) International Application Number: PCT/US94/03862 (22) International Filing Date: 8 April 1994 (08.04.94) (30) Priority Data: 059,215 7 May 1993 (07.05.93) US (71) Applicant (for all designated States except US): APPLE COMPUTER, INC. [US/US]; 20525 Mariani Avenue, Cupertino, CA 95014 (US). (72) Inventor; and (75) Inventor/Applicant (for US only): DAVIDIAN, Gary, G. [US/US]; 247 Sierra Vista Avenue, Mountain View, CA 94043 (US). (74) Agent: FLIESLER, Martin, C.; Fliesler, Dubb, Meyer and Lovejoy, Suite 400, Four Embarcadero Center, San Francisco, CA 94111-4156 (US).</p>	<p>(81) Designated States: AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, ES, FI, GB, HU, JP, KP, KR, KZ, LK, LU, LV, MG, MN, MW, NL, NO, NZ, PL, PT, RO, RU, SD, SE, SK, UA, US, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).</p> <p>Published With international search report.</p>	

(54) Title: METHOD FOR DECODING SEQUENCES OF GUEST INSTRUCTIONS FOR A HOST COMPUTER

(57) Abstract

Emulator performance can be improved by recognizing repeated sequences of the same instruction, or commonly groups of instructions. For example, it is very common to see a three instruction sequence of MOVEM, UNLK A6, and RTS instructions for a 68020 processor in procedure exit code. By looking for these sequences, and combining the operations performed by the separate sequences, overhead of decoding and dispatching the individual instructions in the sequence can be eliminated, and performance improved. Common instruction sequences or repeated sequences in a guest program are detected during emulation of the guest program on a host processor, and performance of the emulation optimized based on the detected sequences. Thus, the emulation logic comprising host instructions embedded within a particular emulation program for a particular guest instruction, detects a particular sequence of guest instructions and in response to detection of the particular sequence bypasses the dispatch logic for guest instructions within the particular sequence. The sequences detected can comprise repeated guest instructions, or common sequences of two or more than two guest instructions.



FOR THE PURPOSES OF INFORMATION ONLY

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

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgystan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

**METHOD FOR DECODING SEQUENCES OF GUEST INSTRUCTIONS
FOR A HOST COMPUTER**

BACKGROUND OF THE INVENTION

5 Field of the Invention

 The present invention relates to the emulation of software written for a given computer on a different computer which executes a different set of instructions; and more particularly to a system for decoding guest
10 instructions into host instructions by the host computer.

Description of the Related Art

 Central processing units for computers are designed to execute a specific set of instructions unique to a
15 particular central processing unit. Thus, a microprocessor in one family, such as the Motorola 68000 family, executes software written in a language unique to the 68000 family, while processors in the Intel 80286 family execute software written with another language
20 which is unique to that family of processors. A need often arises to execute software written for a particular processor in a host processor that utilizes a different language. For the purposes of this application, the language for the host CPU is based on "host
25 instructions", while the language for other CPUs are referred to as "guest instructions".

 Because of the large body of software written for existing processors, such as the Motorola 68000 series, new processors which are designed often attempt to
30 emulate the 68000 series processors in software. This emulation process involves first decoding the 68000 series guest instructions into a sequence of host instructions which accomplish the result intended by the guest instruction. The routines needed to emulate a

given instruction are stored in a host addressable table. For instance, in one prior art system, each guest instruction is used to generate a jump table pointer which points to a table that includes one entry for each
5 of the possible combinations of operational code and addressing mode for a guest instruction. The jump table stores a pointer to the particular code segment adapted to a particular code combination. See, ARRANGEMENT FOR SOFTWARE EMULATION, International Application No.
10 PCT/GB87/00202; invented by MacGregor.

A disadvantage of this prior art technique arises because of the delay involved in composing a jump table pointer based on the guest instruction, looking up a second pointer in the jump table, and then accessing the
15 emulation segment. Because of data access latencies and the like, this can significantly slow down the emulation routine.

At least one prior art system based upon a jump table includes an ability to detect common opcode
20 sequences, so that optimized emulation programs may be used for such sequences. This prior art system is based on eight bit opcodes, and forms a 16 bit index using the two 8 bit opcodes in the sequence. Using the 16 bit index, a jump table entry can be found to an optimized
25 emulation program. This approach works well with eight bit opcodes, which result in an index of manageable size. However for current 16 bit opcode systems, this approach would require a 32 bit index - or with over 4 billion possible entries. The prior art has not provided a
30 practical technique for optimizing the emulation of commonly occurring guest instruction sequences for longer opcodes.

Accordingly, it is desirable to provide an emulation system which enhances the performance of the host

processor in emulation of guest instructions sequences particularly suited to longer opcodes of current processor architectures.

5

SUMMARY OF THE INVENTION

Emulator performance can be improved by recognizing repeated sequences of the same instruction, or commonly used groups of instructions. For example, it is very common to see a three instruction sequence of MOVEM, UNLK
10 A6, and RTS instructions for a 68020 processor in procedure exit code. By looking for these sequences, and combining the operations performed by the separate sequences, overhead of decoding and dispatching the individual instructions in the sequence can be
15 eliminated, and performance improved. The present invention provides a technique in which common instruction sequences or repeated sequences in a guest program can be detected during emulation of the guest program on a host processor, and performance of the
20 emulation optimized based on the detected sequences. According to the present invention, this detection is performed in the execution of the emulation program for the first instruction in the sequence. This approach does add some overhead to the general case of the first
25 instruction when it is not followed by the expected second instruction. However, these optimizations improve performance if the probability of this sequence occurring is very high in the guest program being emulated.

Accordingly, the present invention can be
30 characterized as a system for decoding guest instructions in a host processor. The system includes a guest program store in host addressable memory to store at least one program of guest instructions. An emulation program store in the host addressable memory includes a set of

emulation programs which comprise host instruction routines for respective particular guest instructions. Further, dispatch logic is coupled with the guest program store and the emulation program store to dispatch
5 emulation programs in response to guest instructions in a guest instruction program, and in response to the emulation programs in the emulation program store. Emulation logic comprising host instructions embedded within a particular emulation program for a particular
10 guest instruction, detects a particular sequence of guest instructions and in response to detection of the particular sequence bypasses the dispatch logic for guest instructions within the particular sequence. The sequences detected can comprise repeated guest
15 instructions, or common sequences of two or more than two guest instructions.

According to a further aspect of the invention, both the dispatch logic and the emulation logic which detects guest instruction sequences are based on host
20 instructions embedded within the emulation programs.

Further, the decoding system is further optimized according to the present invention by dividing the emulation program store into a dispatch store and an emulation routine store. The dispatch store is placed in
25 host addressable memory, and stores a set of dispatch entries. Each dispatch entry includes a plurality of host instructions of the emulation program corresponding to a particular guest instruction. The emulation routine store is placed in host addressable memory and includes
30 a set of emulation entries beginning at corresponding emulation table addresses. Each emulation entry in the set includes a host instruction routine for the particular emulation program. The plurality of host instructions in the dispatch entries include a host jump

instruction which causes a jump upon execution by the host processor to an emulation table address of a corresponding emulation entry in the emulation routine store. The emulation entries include host instructions which upon execution by the host processor form an emulation program address to a dispatch entry in response to a next guest instruction, and jump directly to the dispatch table address. Thus, according to this optimization, emulation programs in the emulation program store comprise a dispatch table entry and a host instruction routine in the emulation routine store. Further, the dispatch logic and the logic for detecting guest instruction sequences are embedded in the host instruction routine stored in the emulation routine store.

According to this embodiment, the host system includes a guest instruction pointer store for a guest instruction pointer indicating a guest instruction address, a prefetched guest instruction store for a guest instruction read from the program of guest instructions in response to the guest instruction pointer, and an emulation program pointer store for an emulation program address formed in response to the guest instruction read from the prefetched guest instruction store. Also, the emulation program includes a first macro which upon execution by the host processor forms an emulation program address in the emulation program pointer store in response to the guest instruction in the prefetched guest instruction store. A second macro of host instructions reads the next guest instruction from an address indicated by the guest instruction pointer into the prefetched guest instruction store. A third macro of host instructions causes a jump to the emulation program indicated by the emulation program address in the

emulation program pointer store. A fourth macro embedded within the emulation routine bypasses at least the third dispatch segment upon detection of the particular sequence of guest instruction.

5 Accordingly, a decoding technique for emulating a guest program of instructions in a host processor is provided. The system includes an optimized technique for dispatching emulation programs in response to guest instructions, in combination with the ability to detect
10 common sequences of guest instructions. Upon detection of common sequences, optimized emulation programs can be executed which bypass the dispatching overhead, and which, where appropriate, optimize the emulation of the instruction sequence. The improved emulator optimizes
15 common instruction sequences in the guest programs being emulated, even for guest instructions which comprise long opcodes (e.g., 16 bits).

Other aspects and advantages of the present invention can be seen upon review of the figures, the
20 detailed description, and the claims which follow.

BRIEF DESCRIPTION OF THE FIGURES

Fig. 1 is a schematic block diagram of a computer system implementing the present invention.

25 Fig. 2 is a diagram of a dispatch store according to the present invention.

Fig. 3 is a diagram of an emulation routine store according to the present invention.

30 Fig. 4 is a flow chart illustrating the decoding method according to the present invention.

Fig. 5 illustrates an emulation program according to the present invention.

Fig. 6 illustrates an alternative emulation program according to the present invention.

Fig. 7 is a flow chart of an emulation program which is optimized for a three opcode sequence, which may be common in guest programs.

5 Fig. 8 illustrates an emulation program according to the present invention for a guest instruction which may be a first instruction of a three instruction sequence of guest instructions.

10 Fig. 9 illustrates an emulation program according to the present invention for a guest instruction, which may be the first guest instruction in a two opcode sequence.

Fig. 10 illustrates an emulation program according to the present invention for a guest instruction, which may be repeated.

15 DESCRIPTION OF THE PREFERRED EMBODIMENTS

A detailed description of preferred embodiments of the present invention is provided with respect to Figs. 1-10.

20 I. Emulator System

Fig. 1 illustrates a host processor which is adapted to emulate guest instructions according to the present invention. The host processor includes a host CPU 10 which executes host instructions. Coupled with the host CPU 10 are a set of general purpose registers 11, and a set of special purpose registers 12, implemented as commonly known in the industry as part of an integrated circuit microprocessor, incorporating the CPU 10.

25 The host CPU 10 is coupled to a system bus 13. The bus is also coupled to input devices such as a keyboard and mouse 14, a display system 15, and a memory system 16 such as a disk drive.

The host processor system also includes host addressable memory, which includes an emulation program

store 17, a sequence of guest instructions in a guest program store 18, and other memory 19. The emulation program store 17, according to a preferred embodiment of the present invention, includes an emulation dispatch store 20, and a set of emulation routines 21.

Fig. 2 illustrates the structure of a preferred embodiment of the emulation dispatch store 20. The emulation dispatch store 20 includes a set of instruction pairs, e.g., pair 30, pair 31, pair 32. Each pair includes a host non-jump instruction and a host jump instruction.

According to the preferred embodiment, the dispatch store is an indexed table of 65536 pairs of host instructions, which correspond to the entry points for emulation programs to emulate each of the 65536 possible guest instruction encodings for a Motorola 68020 microprocessor assembly language. The first instruction of the pair will generally perform some operation related to a source operand fetching or addressing. The second instruction of the pair is generally a branch instruction to an emulation routine which resides outside of the dispatch store in the emulation routine store 21. Since each pair of instructions will occupy 8 bytes, the total size of the dispatch store is 512K bytes. The dispatch store may be located, for instance, at addresses (hex)68000000 through (hex)6807FFFF.

The alignment of the beginning of the store is important. In a preferred system, the store starts at either the beginning of a 2 Megabyte address boundary, or a 512 Kilobyte address boundary past the 2 Megabyte boundary. By having this 512K byte block aligned onto a 512K byte address, block address translation registers can be used to perform address translation of a fairly

randomly accessed block of code and eliminate potential thrashing in translation lookaside buffers.

The 512K byte alignment also allows a single instruction to index the dispatch store using a sixteen
5 bit 68020 opcode multiplied by 8. Thus, a single host instruction can shift a 68020 opcode left by 3 bits (multiplied by 8), and insert it into the address of the base of the table to form the address of a dispatch table entry for that opcode.

10 By having this table start in the first 1 Megabyte of a 2 Megabyte aligned boundary, where the second 1 Megabyte of addresses will cause an exception if accessed, it is possible to use an additional address bit to assist in the detection of address errors as described
15 below.

As illustrated in Fig. 2, the host CPU 10 executes a host instruction in order to dispatch a guest instruction by accessing the first instruction in an instruction pair in the dispatch store. The first
20 instruction executes, and then a second instruction in the dispatch pair is executed. This second instruction includes a jump instruction which goes to an emulation block in the emulation routine store 21.

Fig. 3 illustrates the implementation of the
25 emulation routine store 21. The emulation routines are allocated to a 64K block of bytes for host instruction routines to which the jump instruction in the dispatch entry branches. In general, the first two host instructions in an emulation program reside in an entry
30 in the dispatch store, while the remaining instructions, which are referred to as emulation routines, reside in the emulation routine block of code. This block may be located, for instance, at addresses (hex) 68080000 through (hex) 6808FFFF.

As above, the alignment of the block of code for the emulation routines is also important. In a preferred system, it needs to start at the beginning with 64K byte boundary.

5 During emulation, it is frequently desirable to compute the address of a guest instruction within the emulated block, such as by computation of a PC relative address. The host architecture in a preferred system may not provide an instruction to compute a PC relative
10 address. By storing the address of the beginning of the emulation routine block so that it has zeros in the 16 least significant bits in a host register, referred to as `code_ptr`, a computation of the address of any label within this 64K byte block of code can be optimized by using the
15 value in `code_ptr` as a code base by doing an OR immediate, with a 16 bit immediate value as the offset.

 Within the 64K byte block of emulation routine code, there is additional attention paid to code alignment. In particular, the emulation blocks are aligned into blocks
20 which match the processor caching routing used to retrieve the code. In a preferred system, a processor cache uses 32 byte cache blocks in a cache line of 64 bytes, and the emulation blocks are packed into aligned 32 and 64 byte blocks.

25 Thus, as illustrated in Fig. 3, the emulation routine store 21 may include a plurality of emulation routines including emulation block 40, emulation block 41, emulation block 42, emulation block 43, emulation block 44, etc. Each of these blocks 40-44 is either a 32
30 or 64 byte block.

 A particular emulation block, e.g., block 42, is entered by a jump from the dispatch table, and ends with an instruction to dispatch the next guest instruction.

As illustrated in Fig. 3, some emulation blocks may include effective address calculation routines, such as block 43. Such effective address routines are entered by a jump from the dispatch table as described below, and
5 end with a jump to a return address of an emulation block within the emulation routine memory.

Fig. 4 illustrates the emulation decoding process according to the present invention. As mentioned above, the host processor 10 shown in Fig. 1 includes a
10 plurality of special purpose and general purpose registers. The general purpose register `disp_table` stores a pointer to a dispatch table entry. Also, at this point in the decoding sequence, a special purpose register labelled `ctr` will contain the same value as `disp_table`. A
15 general purpose register labelled `pc` stores a guest instruction pointer 2 bytes past a current guest instruction. A general purpose register labelled `prefetch_data` stores a next guest instruction, as indicated at block 50. As mentioned above, the emulation routines
20 include host instructions distributed within the routines to carry out the decoding process. Thus, each emulation program will include a DISPATCH macro, generally 51, which does inter-instruction processes which may be required between guest instructions, and causes a jump to
25 the dispatch table entry indicated by the pointer in register `ctr`. The emulation program also includes a macro referred to as DECODE1 macro 52, which takes the next guest instruction from the `prefetch_data` register, multiplies that instruction by 8, and inserts the results
30 in the general purpose register labelled `disp_table` so that it forms an address of a dispatch table entry.

A next macro 53 within an emulation program referred to as DECODE2 macro, copies the value in the `disp_table` register to the special purpose register `ctr`.

5 A final macro referred to as PREFETCH macro 54 is included within an emulation program. The PREFETCH macro 54 advances the guest instruction pointer in register `pc` by 2 bytes, then causes a prefetch of the next guest instruction from the address indicated by the pointer in register `pc` and places the prefetched instruction in the
10 general purpose register `prefetch_data`. The final macro in a given emulation routine is the DISPATCH macro 51. Thus, as illustrated in Fig. 4, an emulation program for a particular guest instruction begins immediately after the jump instruction of the DISPATCH macro 51.

15 The structures of alternative emulation programs, according to the present invention, are shown in Figs. 5 and 6. Fig. 5 illustrates the general case. Line 1 is the first instruction `INST1` stored in a dispatch entry in the dispatch store. Generally, this is a non-jump
20 instruction relevant to addressing mode of the guest instruction. Line 2 stores the second instruction `INST2` stored in the dispatch entry. Generally, this is a jump to a block of code in the emulation routine store.

25 Line 3 corresponds to the instruction or instructions at the beginning of an emulation block. Within an emulation program, the DECODE1 macro occurs next, as indicated at line 4. After the DECODE1 macro in line 4, an instruction or instructions may be included relevant to emulation of the guest instruction. Next, a
30 DECODE2 macro is executed as indicated at line 6. The DECODE2 macro in line 6 may be followed by other instructions indicated by line 7, relevant to the decoded guest instruction. Next in the sequence, a PREFETCH macro of line 8 is executed. The PREFETCH macro may be

followed by an instruction or instructions represented by line 9 of the emulation program. Finally, the DISPATCH macro is executed as indicated at line 10. The first instruction of the DISPATCH macro is a conditional jump
5 to the value in the ctr register, if no special conditions are pending.

Fig. 5 illustrates the distribution of the DECODE1, DECODE2, PREFETCH, and DISPATCH macros within an emulation program. It will be appreciated by those of
10 skill in the art that the presence of instructions between such macros may or may not occur. However, by distributing the macros among the instructions of the emulation program, the programmer can take advantage of any data or instruction access latency occurring in the
15 program to improve performance.

Fig. 6 illustrates an emulation program which is used when a single instruction in a dispatch entry in the dispatch table is insufficient to handle addressing mode issues. In this case, the first instruction, INST of the
20 dispatch entry, is shown in line 1. This instruction is a non-jump instruction which sets the value of a return address in a host register labelled rtn_addr which is a general purpose register in a preferred embodiment of the present invention. The next instruction in line 2 of
25 Fig. 6 is the instruction INST2 in the dispatch entry. This instruction causes a jump to an effective address routine stored in the emulation routine store, as illustrated at element 43 of Fig. 3.

Line 3 of Fig. 6 illustrates the beginning
30 instructions of the effective address routine. Line 4 of the program shown in Fig. 6 is an instruction which results in moving the return address from the register rtn_addr to a special purpose register lr, which is used for jump addresses by the host processor. This may or may

not be necessary in a given implementation of the present invention. Line 5 of the program illustrates the final instruction of the effective address routine, which requires a jump to the return address stored in the register lr. This jump results in execution of the instructions illustrated at line 6, starting the emulation block for the guest instruction. This emulation block will include the DECODE1 macro, line 7, possibly other instructions, line 8, the DECODE2 macro, line 9, the PREFETCH macro, line 10, possibly other instructions, line 11, and the DISPATCH macro, line 12, as described above with respect to Fig. 5.

II. Multi-Instruction Sequence Emulation

The emulation system according to the present invention may be further optimized for sequences of guest instructions which are expected to be relatively common in the guest code to be emulated. For instance, the move multiple opcode MOVEM in the 68020 architecture may be a first instruction in a three instruction sequence, including MOVEM, unlink UNLK, and return-to-subroutine RTS, which commonly occurs in many programs written for the 68020 processor. Thus, Fig. 7 illustrates how an emulation program for the MOVEM opcode may be implemented according to the present invention.

The MOVEM emulation program, after dispatching, as described above, will include a section which begins execution of the opcode, including the DECODE1, DECODE2, and PREFETCH macros (block 100). After the prefetching, the next opcode (2nd) can be tested to determine whether it is the UNLK opcode (block 101). If it is not, then the expected sequence is not occurring, and the MOVEM instruction is completed (block 102).

If at block 101, the emulation program detects a UNLK instruction, then the algorithm tests whether special conditions are pending, such as an interrupt, or instruction tracing condition (block 103). If a special
5 condition is pending, then the algorithm branches back to complete the MOVEM instruction at block 102, because the special condition must be handled between guest instructions.

If no special condition is pending at block 103,
10 then the next opcode (3rd) is tested (block 104). If the third opcode is not the RTS instruction, then the predicted three instruction sequence is not found, and the algorithm branches to complete a combined MOVEM and UNLK instruction program (block 105). If the next opcode
15 is found to be RTS in block 104, then the algorithm branches to complete the combined MOVEM, UNLK, and RTS instruction sequence (block 106).

Thus, it can be seen that for the combined sequence, the overhead of the decoding and dispatching logic for
20 UNLK and RTS is bypassed.

Fig. 8 illustrates an emulation program for the MOVEM instruction, or a similar instruction which may be the first instruction in a three guest instruction sequence. As illustrated in Fig. 8, the first guest
25 instruction is dispatched by the emulation program of the previous instruction, as described above. Thus, the first guest instruction emulation program includes instruction 1 on line 1 and instruction 2 on line 2 which are found in the dispatch table. Instruction 2 on line
30 2 causes a jump to line 3 of the program, found in the emulation program store. These instructions for the MOVEM instruction will need a MOVEM opcode extension. This opcode extension will have been loaded by the previous emulation program in the `prefetch_data` register, as

it consists of the next two bytes in the guest instruction sequence. Thus, an additional prefetch is needed to be executed to find the next guest instruction. So, the pc value is then incremented by two bytes and a
5 prefetch operation to fetch the instruction from the address pointed to by the register pc to a general purpose register gpr, other than the prefetch_data register is executed (line 4). The emulation program tests whether the second guest instruction, which is now indicated by
10 the value in the general purpose register is the expected second guest instruction of the sequence and saves the result (line 5).

Line 6 of the program indicates that the extension data is used by the emulation program. An additional
15 instruction or group of instructions may be executed (line 7). Lines 8 and 9 of the program illustrate that the DECODE1 and DECODE2 macros are executed. Line 10 indicates that an instruction or instructions may be interspersed between the DECODE2 macro and the PREFETCH
20 macro on line 11. After the DECODE1, DECODE2 and PREFETCH macros are executed, a check for special conditions, such as an interrupt or instruction tracing mode, is made (line 12). As indicated at line 13, the program will branch if the sequence had been detected in
25 line 5, and no special conditions were found. If the sequence is not detected, then instructions indicated at line 14 are executed to complete emulation of the MOVEM instruction. Finally, the DISPATCH macro is executed to jump to the dispatch table entry for the second guest
30 instruction (line 15).

In Fig. 8, if the branch is taken at line 13, then the program moves to line 16. At this point, instructions may be executed. Line 17 of Fig. 8 shows that the program then tests for the third expected guest

instruction in the sequence by looking at the value in the `prefetch_data` register. As indicated at line 18, the algorithm will branch if the sequence is detected. If not, the instructions indicated by line 19 are executed to carry out a combined two instruction sequence. The `DECODE1` and `DECODE2` macros are executed, as indicated at lines 20 and 21. Line 22 indicates the possibility of interspersed instructions for the emulation program between the macros. At line 23, the `PREFETCH` macro is executed to prefetch the fourth guest instruction in the sequence. Next, the `DISPATCH` macro is executed to jump to the dispatch table entry for the third guest instruction in the sequence (line 24).

If at line 18 of the program the branch was taken, then for the program illustrated in Fig. 8, line 25 is executed. This line indicates that instructions in the emulation program are executed. Line 26 illustrates that an additional fetch operation is executed to retrieve the fourth guest instruction into the `prefetch_data` register. For the `RTS` instruction in this sequence, the `PREFETCH` macro is replaced by instructions to retrieve the target of the return from subroutine instruction into the `prefetch_data` register.

Lines 27 and 28 correspond to the standard `DECODE1` and `DECODE2` macros in the emulation system. Line 29 indicates the presence of interspersed instructions to complete the combined three guest instruction emulation. Line 30 illustrates the presence of the `PREFETCH` macro to retrieve the fifth guest instruction in the sequence. Line 31 ends up the routine with a `DISPATCH` macro which causes a jump to the dispatch table entry of the fourth guest instruction in the sequence.

Thus, it can be seen that the emulation program includes logic which bypasses the dispatching of

instructions in the detected sequence. This greatly reduces the overhead involved in executing common guest instruction sequences.

5 Fig. 9 illustrates an emulation program for a sequence of guest instructions, such as a repeated sequence. As can be seen in Fig. 9, the first guest instruction is dispatched by the previous emulation program, and instructions 1 and 2 from the dispatch table entry are executed (lines 1 and 2). Instruction 2 causes
10 a jump to the emulation program store and execution of the sequence beginning with line 3.

Line 3 in Fig. 9 illustrates that the emulation program will include instructions that test the `prefetch_data` register for the sequence. This register will
15 have been filled with the second guest instruction in the sequence by the `PREFETCH` macro of the previous guest instruction.

As indicated on line 4, the algorithm branches if the sequence is detected to line 11. If the sequence is
20 not detected, then the algorithm continues with instructions indicated at line 5 to complete the emulation program. The emulation program would also include the `DECODE1` and `DECODE2` macros as indicated at lines 6 and 7. Line 8 indicates that instructions of the
25 emulation program may be interspersed with the macros involved in decoding and prefetching instructions. At line 9, the `PREFETCH` macro is executed to retrieve the third guest instruction in the sequence. Line 10 is the `DISPATCH` macro which jumps to the dispatch table entry
30 for the second guest instruction. If at line 4, the branch was taken, a `PREFETCH` macro is executed to retrieve the third guest instruction. This macro is necessary because the `DECODE1`, `DECODE2`, and `PREFETCH` macros of lines 6, 7, and 9 were bypassed by the

detection of the sequence of guest instructions. Then, the instructions at line 12 are executed for the combined emulation program. Next, the emulation program will include a DECODE1 macro as indicated at line 13. Line 14
5 indicates the possible distribution of instructions within the emulation program. Line 15 indicates the DECODE2 macro which loads the ctr register with the address of the dispatch entry for the third guest instruction. Line 16 indicates the PREFETCH macro for
10 retrieving the fourth guest instruction in the sequence. Line 17 is the DISPATCH macro which causes a jump to the dispatch table entry for the third guest instruction.

Thus, Fig. 9 illustrates the bypassing of decode and dispatch logic in the emulation program for a common
15 opcode sequence, such as a repeated pair of opcodes.

Fig. 10 illustrates a program which is optimized for decoding repeated sequences of instructions. In this example, the first instruction on line 1 is a test for repeat and the second instruction is a jump from the
20 dispatch table entry to an emulation routine on line 3 as shown in Fig. 10. Thus, lines 1 and 2 of Fig. 10 correspond to the dispatch table entry for the first guest instruction.

At line 3, the routine executes an instruction or
25 instructions relevant to the guest instruction being decoded. Next, the DECODE1 macro is executed, followed by the PREFETCH macro for the third guest instruction (lines 4 and 5). After the PREFETCH macro, an instruction or instructions for emulating the guest
30 instruction are executed (line 6). In line 7, the algorithm then branches if a repeat had been detected to line 11 of the routine. If no repeat had been detected, then the DECODE2 macro is executed, as indicated at line 8. This macro is followed by instructions indicated at

line 9 which wrap up the emulated guest instruction. Line 10 indicates the DISPATCH macro is executed. This results in dispatching of the (2+N)th guest instruction, where "N" is the number of times that a repeated instruction had been detected. Thus, if no repeat is
5 detected at line 7, then the second guest instruction is dispatched at line 10.

If the branch had been taken at line 7, then the algorithm goes to line 11 to test for a repeat once
10 again. Thus, the second and third guest instructions can be compared to determine whether a repeat has occurred at this point, because the PREFETCH macro on line 5 had prefetched the third guest.

In line 12 of the routine, a branch is taken if a
15 special condition is detected to line 8 to complete the execution of the current guest instruction. In line 13, additional instructions are executed to handle the combined execution of repeated instructions. At line 14, the DECODE1 macro is executed, followed by the PREFETCH
20 macro in line 15. The PREFETCH macro on line 15 prefetches the (3+N)th guest instruction, where "N" again is the value indicating the number of times that a repeat had been detected.

At line 16, instructions are executed relevant to
25 emulation of the guest instruction sequence. At line 17, the algorithm branches if a repeat had been detected at line 11 back to line 11. The algorithm continues in this loop from lines 11 through 17, until the repeated sequence ends. Line 18 of the routine causes a branch to
30 line 8 to wrap up the sequence if the branch at line 17 is not taken.

III. Details of a 68020 Emulation on a Power Architecture

The present invention may be further understood with reference to or detailed information concerning emulation of the Motorola 68020 microprocessor guest code on an IBM
 5 POWER microprocessor architecture. Thus, the internal design of a 68020 emulator for the POWER architecture processor is provided below.

POWER Registers

10 The POWER architecture defines 32 general purpose 32 bit registers (actually they can be 64 bits, but the emulator just uses the 32 bit architecture) referred to as r0...r31. There are 32 double precision 64 bit floating point registers referred to as f0...f31, which are not used
 15 at all by the emulator. There are 4 additional special purpose 32 bit registers used by the emulator, they are called cr (condition register), xer (exception register), ctr (counter), and lr (link register).

In the source code, the general purpose register are referred to by names, instead of r0..r31. These name will
 20 be used in the remainder of the document. The current assignments are as follows, although they can easily be rearranged.

25	r0	zero
	r1	a7
	r2	(unused)
	r3	addr
	r4	data
30	r5	rtn_addr
	r6	immed_data
	r7	base_disp (also called scaled_index)

	r8...r15	d0...d7
	r16...r22	a0...a6
	r23	(unused)
	r24	pc
5	r25	sr_and_flags
	r26	ccr_x
	r27	prefetch_data
	r28	vbr
	r29	disp_table
10	r30	code_ptr
	r31	EmulatorStatePtr

The zero register contains a constant value of zero, it is never changed. Assigning this to register r0 is also convenient due to the POWER architecture base address quirk that does not allow r0 to be used as a memory base register.

The register d0...d7, a0...a7, pc, sr_and_flags, ccr_x, and vbr are used to hold corresponding 68020 register state. This is described in detail later in this document.

Registers addr, data, rtn_addr, immed_data, and base_disp are five temporary scratch registers used during the emulation of a 68020 instruction. Although they can be used for many different purposes, their names describe how they are used by effective address calculation routines, which are used during the emulation of many of the opcodes.

The prefetch_data register generally contains the sign extended 16 bit data value pointed to by the pc register.

The disp_table register points to an entry in the 512KB instruction dispatch table. The opcode being dispatched

to is inserted into this register to index the table, this is described in more detail later.

The `code_ptr` register points to the beginning of the 64KB block of code that contains the *emulation routines*.
5 this block is 64KB aligned so that a 16 bit immediate value can be "or"ed with this register to point to any address with this 64KB block of code.

`EmulatorStatePtr` points to the base of the memory area that is used to store less frequently used emulator state that cannot be contained in the POWER registers.
10

The POWER special purpose `lr` register is available for use during the emulation of a 68020 instruction, and does not correspond to any 68020 register state.

The POWER special purpose `ctr` register is available for use during the emulation of a 68020 instruction, and does not correspond to any 68020 register state. It is used by convention to hold the address of the first POWER instruction to be executed to emulate the next 68020 instruction.
15

The POWER special purpose `xer` register is used to hold the V and C bits of the 68020 CCR register. It also contains the POWER SO bit, as well as the byte count to be used by POWER string instructions.
20

The POWER condition register `cr` is used to hold the N and Z bits of the 68020 CCR register. The low 16 bits are available for general use during the emulation of a 68020 instruction. The 4 bit condition register fields `cr1` and `cr2` are not used, the 4 bits of `cr3` are used for global flags related to interrupts and special conditions which are described later.
25
30

68020 Register State Assignments

The 68020 has a number of general purpose and special purpose registers, some of which are only accessible in supervisor mode. All of these registers contents must be maintained by the emulator. In some cases, the bits in these registers may be distributed in a number of different places within the POWER architecture, but the emulator will gather/scatter the bits whenever it encounters a 68020 instruction that accesses the entire register. In other cases, there may be multiple copies of a register contents. Many of the special purpose registers are stored in memory pointed to by EmulatorStatePtr. The 68020 registers and their POWER locations are as follows.

15	D0...D7	d0...d7	
	A0...A6	a0...a6	
	A7	a7	(currently active stack pointer)
	PC	pc	(POWER version does not always point to executing instruction)
20	PC	trace_pc(EmulatorStatePtr)	(valid when tracing enabled)
	CCR	cr/xer/ccr_x	(bits are distributed)
	SR	sr_and_flags	(upper byte of SR only)
25	USP	saved_usp(EmulatorStatePtr)	(a7, when usp is active stack)
	ISP	saved_isp(EmulatorStatePtr)	(a7, when isp is active stack)
	MSP	saved_msp(EmulatorStatePtr)	(a7, when msp is active stack)
30	VBR	saved_vbr(EmulatorStatePtr)	(duplicate copy in vbr)

SFC saved_sfc(EmulatorStatePtr)
DFC saved_dfc(EmulatorStatePtr)
CACR saved_cacr(EmulatorStatePtr)
CAAR saved_caar(EmulatorStatePtr)

5

The 68020 registers d0...d7/a0...a6 are in POWER registers. The 68020 has three stack pointers, and the active stack pointer will be in register a7, while the remaining two inactive stack pointers will reside in memory. The memory copy of the active stack pointer is not used and inconsistent while that stack pointer is selected as the active stack pointer. When the selection of the active stack pointer is changed, the register copy of the old stack pointer will be written to memory, and the new register copy of the active stack pointer will be read from memory. It should be noted that register a7 is assigned to POWER register r1, which is register used for the native POWER stack pointer. This is to allow a single stack model in a mixed emulated and native environment.

10
15
20

In the 68020, the pc generally points to the beginning of the instruction that is currently being executed. During emulation, the pc register advances as the instruction is decoded and executed, and generally points somewhat past the beginning of the instruction being executed. At the beginning of the execution of an instruction, the pc always points two bytes (16 bits) past the beginning of the instruction, which may actually point to the next instruction. Since this offset is constant, it is always possible to compute the actual pc at an instruction boundary. When the 68020 instruction trace mode is active, the exception frame that is generated after the execution of a traced instruction

25
30

needs to contain the pc of the beginning of the instruction that has just completed. Since it is generally not possible to compute the size of the instruction that just completed, or worse yet, it may
5 have been a branch instruction which computed a completely new pc, there is a memory copy called trace_pc. When trace mode is active, the pc of an instruction that is about to execute is save in the memory based trace_pc, so that the starting pc of the instruction can be
10 determined when the instruction completes. Since there is a performance penalty associated with this computation and updating, this is only performed when trace mode is enabled.

The 68020 CCR register consists of five condition
15 code bits, called X, N, Z, V, and C. During emulation, these are treated as five separate bits which are in three different registers, instead of a single field of five bits. The X bit is stored in bit 2 of the POWER register named ccr_x. This bit position corresponds the
20 position of the CA bit in the POWER special purpose XER register. The N bit is stored in bit 0 of the POWER cr register, this corresponds to the LT condition bit of cr0. The Z bit is stored in bit 2 of the cr register, this corresponds to the EQ condition bit of cr0. The V bit is
25 stored in bit 1 of the POWER special purpose XER register, which is the OV flag. The C bit is stored in bit 2 of the XER register, which is the CA flag. Most of the 68020 data movement and logical operations only update four of the five condition codes. They leave the
30 X bit unchanged, set the N and Z bits to indicate if the result is negative or zero, and always clear the V and C bits. Using this arrangement, a single POWER instruction

can move data from one register to another, and update some or all of these four bits of the CCR as follows.

```

ao.  dst,src,zero  ;# move data, update N,Z, clear V,C
or.  dst,src,zero  ;# move data, update N,Z
5   ao  dst,src,zero  ;# move data, clear V,C
    caxo dst,src,zero  ;# move data, clear V
    a   dst,src,zero  ;# move data, clear C

```

Most of the 68020 arithmetic and shift operations
10 update the X bit to the same value as the C bit. Since the C bit is in the XER register, a simple move from the XER register into the ccr_x register is all that is required to update the X bit. It should be noted that the 68020 X and C bits are set to 1 if there is a borrow
15 during subtraction, while the POWER (and most other RISC processors) set the CA bit to 0 if there is a borrow during subtraction. This will require the CA bit to be complemented before saving it as the X and C bits. The same inversion is needed when the X bit is used as a
20 borrow-in for the 68020 SUBX instruction. By using the following instruction pair, it is possible to perform a subtraction followed by an addition, which will set the CA flag to correspond to the 68020 conventions.

```

25   sfo. dst,src,dst  ;# dst <- dst-src, update N,Z,V
    a   tmp,dst,src  ;# update C, ignore result

```

The upper byte of the 68020 SR register is stored in the low 8 bits (24.31) of the sr_and_flags register. The
30 high 16 bits (0..15) of this register contain test-mode enable flags. Bit 20 is the flag_group_1_active bit which

indicates that a group 1 exception (Bus Error or Address Error) is being processed, and is used to detect a *Double Bus Fault* situation. Bit 22 is the *flag_odd_pc* bit, which is used for *Address Error* detection. Bit 23 is the *flag_trace_pending* bit, which indicates that the 68020 instruction currently being executed is being traced, and needs to generate a trace exception when it completes.

Many of the 68020 special purpose registers that are accessed via the MOVEC instruction are stored in memory, because they are infrequently accessed, and there are not enough POWER registers available to hold all of them. An exception to this is the *vbr* register, there is a dedicated POWER register that is used to hold a copy of the *vbr* register contents, however the memory copy is also kept up to date. The various stack pointers are also an exception. Since only one of the three stack pointers can be selected at a time, the register *a7* is used for the selected stack pointer, and the remaining two inactive stack pointers are stored in memory.

20

The Dispatch Table

The dispatch table is an indexed table of 65536 pairs of POWER instructions, which correspond to the entry points for the routines to emulate each of the 65536 possible 68020 instruction encodings. The first instruction of the pair will generally perform some operation related to source operand fetching or addressing. The second instruction of the pair is generally a branch instruction to an *emulation routine* which resides outside of this table. Since each pair of instructions will occupy 8 bytes, the total size of this table is 512K bytes. Currently this table is located at addresses \$68000000..\$6807FFFF, and can reside in Read

30

Only Memory (ROM). The register `disp_table` is used to address this table.

5 The alignment of the beginning of the table is very important, it needs to start at either the beginning of a 2MB boundary, or 512KB past the beginning. By having this 512KB block aligned to a 512KB address, the POWER Block Address Translation (BAT) registers can be used to perform the address translation of this fairly randomly accessed block of code, and eliminate potential thrashing
10 of the TLBs.

The 512KB alignment also allows a single POWER instruction to index the table using the 68020 opcode times 8. The following instruction will shift the opcode left by 3 bits (multiply by 8), and insert it into the
15 address of the base of the table, forming the address of the table entry for that opcode.

```
rlimi disp_table,opcode,3,0x0007FFF8
```

20 By having this table start in the first 1MB of a 2MB aligned boundary, where the second 1MB of addresses will cause an exception if accessed, it is possible to use an additional address bit to assist in the detection of Address Errors (see discussion later in this document).

25 The Emulation Routines

There is a block of 64K bytes allocated for the POWER instructions the Dispatch Table branches to. In general, the first two POWER instructions in the emulation of a 68020 instruction reside in the Dispatch
30 Table, while the remaining instructions, which we refer to as the Emulation routines, reside in this block of code. Currently this block is located at addresses \$68080000..\$6808FFFF, and can reside in Read Only Memory

(ROM). The register `code_ptr` contains the address of the beginning of this block.

Just like the Dispatch Table, the alignment of the block of code for the Emulation routines is also very important, it needs to start at the beginning of a 64KB boundary.

In the emulator, it is frequently desirable to compute the address of an instruction within this block of code. The POWER architecture does not provide an instruction to compute a pc-relative address. The register `code_ptr` points to the beginning of the block, and there is a label in the source code called `cb` which marks the code base. To easily compute the address of any label within this 64KB block of code, the following instruction can be used.

```
ori addr,code_ptr,label-cb
```

Within the 64KB block of code, there is additional attention paid to code alignment. The 601 processor cache has 32 byte cache blocks, and a cache line consisting of 2 cache blocks, or 64 bytes. To improve locality, and reduce the number of bytes of code that needs to be fetched when there is a cache miss, the routines are packed into nicely aligned 32 or 64 byte blocks.

68020 Instruction Prefetching

On the 601 processor, as well as most other RISC processors, there is some latency associated with memory read operations, and attempts to use the results of a load instruction, in the very next instruction will usually result in a pipeline stall. To improve performance, and minimize these stalls, it is very

desirable to issue memory reads several instructions before attempting to use the data that they read.

5 Since the emulator needs to read all of the 68020 opcode and operand bytes in order to emulate an instruction, performance can be improved by issuing these reads long before the data is needed. To accomplish this, the emulator uses a register called `prefetch_data` to read (or pre-fetch) the next 16 bits (sign extended) of the instruction stream into, as soon as the current 16
10 bits have been consumed. The register `pc` is used to point to the position within the 68020 instruction stream that has been read. The POWER architecture provides an efficient instruction that can both advance the `pc` register, and read the data pointed to by the updated `pc`.
15 The instruction is as follows, and there is also a macro called `PREFETCH` that is used within the emulator source code.

```
lhau prefetch_data,2(pc)
```

20 The prefetched data is always read 16 bits at a time, and sign extended, because the 68020 opcodes and extension words are most often organized in 16 bit groups. The sign extension is useful for the addressing modes that use a 16 bit signed displacement that is added
25 to an A-register or the PC register.

68020 Instruction Decoding

By using many of the concepts introduced above, the four basic steps required to decode and emulate a simple
30 68020 instruction can now be described. The four steps are referred to as `DECODE1`, `DECODE2`, `PREFETCH`, and `DISPATCH`. For simplicity, we will assume that this is a 16 bit opcode that does not perform any useful operation.

Since this is a very pipelined sequence of events, and we must start somewhere in the pipeline, we will begin at the first instruction in the dispatch table for this opcode, and after going through the remaining
5 stages, we will see how we get back here after completing the remaining phases.

Upon entry, the following registers are setup as follows. The `disp_table` and `ctr` registers contains the address of the dispatch table entry for this opcode (the
10 POWER address that we are currently executing at). The `pc` register points 2 bytes past the 68020 opcode that we are about to emulate. The `prefetch_data` register contains the sign extended 16 bit value that the `pc` register points to (in this example, it is the next 68020 opcode to be
15 emulated).

The first phase is `DECODE1`, this phase is begins the decoding of the next 68020 instruction to be emulated. In this example, we are assuming that the current 68020 instruction consists of just a 16 bit opcode, and does
20 not have any extension words. If there were extension words, they would need to be consumed by `PREFETCHing`, until `prefetch_data` contains the opcode of the next 68020 instruction, and `pc` points to that opcode. The `DECODE1` operation takes the next 68020 opcode that is in the
25 `prefetch_data` register, multiplies it by 8, and inserts it into the `disp_table` register, forming the address of the dispatch table entry for the next 68020 instruction. This is done in a single POWER instruction as follows, the macro `DECODE1` performs this instruction.

30 `rlimi disp_table,prefetch_data,3,0x0007FFF8`

Since `DECODE1` was the first of the two POWER instructions that reside in the dispatch table entry for

this 68020 instruction, the second instruction must be a branch out of the dispatch table, and into an emulation routine. This is not considered to be one of the phases of the decoding process, but rather a necessity imposed
5 by the two instruction limit with a dispatch table entry. In this example we will assume that this branch is as follows.

b continue

10 The second phase is DECODE2, which in this example will occur in the first POWER instruction of the emulation routine. DECODE2 simply takes the dispatch table entry address that was computed by DECODE1, and moves it into the POWER ctr register. This is because the
15 POWER branch instructions cannot branch to addresses contained in the general purpose registers, and can only branch to addresses in either the ctr or lr special purpose registers. The DECODE2 phase is done in a single POWER instruction as follows, the macro DECODE2 performs this
20 instruction.

continue:

mtctr disp_table

25 The third phase is PREFETCH, which in this example will occur in the second POWER instruction of the emulation routine. As described earlier, PREFETCH will advance the pc register by 2 bytes, and read the sign extended 16 bit value at that address into the prefetch_data register. We need to prefetch at this time, because we
30 have consumed the previous contents of the prefetch_data register, which had contained the 16 bit 68020 opcode for next instruction to be emulated. This will setup the prefetch_data register with the first extension word (if

any) associated with the next opcode, or the opcode of the instruction following the next instruction. As shown earlier, the PREFETCH phase is done in a single POWER instruction as follows, the macro PREFETCH performs this instruction.

```
5      lhau  prefetch_data,2(pc)
```

The fourth and final phase is DISPATCH, which in this example will occur in the third and fourth POWER instructions of the emulation routine. There are two POWER instructions needed for this phase, but in general the second one never gets executed. The DISPATCH phase completes the emulation of the current 68020 instruction, and begins the emulation of the next 68020 instruction. Since this marks the boundary between two 68020 instructions, any special processing that needs to occur between two 68020 instructions must happen here. Instruction Trace exceptions, and 68020 interrupt processing are examples of special events that need to be processed on instruction boundaries. There is a bit in the POWER cr register referred to as cr_special_event_pending, which gets set whenever any of this special handling is needed. How this bit gets set will be described later, but for now, lets just assume that it is cleared. Since the dispatch table entry address for the next 68020 instruction is already loaded into the POWER ctr register, the DISPATCH phase simply needs to branch to this address when there are no special events pending, or branch to a common routine to process pending special events. This final phase is done in two POWER instructions as follows, the macro DISPATCH performs these instructions.

```
25      bfc   cr_special_event_pending
30      b     process_special_event
```

By breaking the decoding and dispatching process into these simple phases, the instructions to perform the various phases can be distributed between other instructions within the emulation routines to execute in gaps where the processor would have otherwise stalled due to memory access latency.

Effective Address Computation

The 68020 architecture has a number of different addressing modes, some are very simple, and some can become very complex. Since the effective address, and the data that it points to, is generally needed very early in the emulation of a 68020 instruction, it is convenient to have a number of Effective Address Computation routines that can run be selected based upon the addressing mode, and a common emulation routine to implement the operation independent of the addressing mode used.

In some cases, the entire Effective Address Computation can occur in a single POWER instruction, and can be placed in the Dispatch Table. In other cases a short (or possibly long) subroutine is needed. Since there is only room for two instructions in the dispatch table, one method of performing a subroutine call would be to have the first instruction contain a call to the subroutine, which would return to the second instruction, which would be a branch to the emulation routine. This would result in a branch to a branch instruction, on the initial dispatch, and a branch to a branch instruction when returning from the Effective Address subroutine. This type of branching does not perform well on the POWER processor. It is more desirable to have the Effective Address calculation subroutine return directly to the @ first instruction of the emulation routine.

To reduce the number of branches, a slightly different subroutine calling convention is used. The first instruction in the dispatch table will load the address of the emulation routine into the register
5 `rtn_addr`, and the second instruction will branch to the Effective Address subroutine. The subroutine will move the address of the emulation routine from `rtn_addr` into the `lr`, and when it returns, it will return to the emulation routine.

10 routine, the `code_ptr` register is used as a base address. An example of what the two instructions in the dispatch table may look like is as follows.

```
ori   rtn_addr,code_ptr,not_l_mem-cb
b     cea_l_30
```

15

There are some register conventions used by the effective address routines. As mentioned before `rtn_addr` is used to pass the return address, and may be used as a scratch register by the Effective Address routine. The register `addr` is used to return the effective address (if one is computed). The register `data` will contain the data that was read from the effective address (if it was a Fetch Effective Address routine), or will remain unchanged. The register `immed_data` is used to return the
20 immediate operand, or opcode extension word that follows to opcode but precedes the Effective Address extension words, for Immediate/Extended Effective Address routines. The register `base_disp` is used as a scratch register.

25 The 68020 modes 6n and 73 indexed addressing modes can be very complex. In it's simplest form, any of the sixteen 68020 A/D registers can be used as an index. The index can be treated as a sign extended 16 bit quantity, or a full 32 bit value, and can be multiplied by 1, 2, 4,
30

or 8, added to a base register, and added to a sign
extended 8 bit displacement. Most of the information
needed to compute the address is in a secondary extension
word. There is an additional bit that indicates that
5 even more complex addressing options are available. To
quickly decode all of these options, there is a 256 entry
table of instruction pairs that is at the beginning of
the emulation routine code block. The `code_ptr` register
points to the base of this indexed addressing mode decode
10 table.

Address Error Processing

Since all instructions in the 68020 architecture are
a multiple of two bytes long, the 68K does not allow
15 instructions to begin on odd addresses. If a branch to
an odd address is attempted, an Address Error exception
is generated. Since this exception is an indication of
a programming error, and is not a normal occurrence, the
emulator would like to spend as little time as possible
20 checking for this condition. There have been two methods
used in the emulator to detect branches to odd addresses,
each has different performance characteristics.

The first method (which is no longer used) consists
of two phases, called `CHECK_ODD_PC1` and `CHECK_ODD_PC2`.
25 Each phase consisted of a single instruction. The first
instruction would move the low bit of the `pc` register into
bit 31 of the `cr` register. The second instruction would
"or" bit 31 into the special event pending flag. This
would cause the `process_special_event` routine to be entered
30 if the new `pc` value was odd. That routine would check to
see if the `pc` was odd, and generate an Address Error
exception if it was. The `POWER` instructions used by the
two phases are as follows.

```
mtcrf 0x01,pc
cror cr_special_event_pending,cr_special_event_pending,31
```

5 The second method consists of a single phase, called
CHECK_ODD_PC1, which is a single POWER instruction. This
instruction will insert the low bit of the pc register
into bit 11 of the disp_table register. Assuming that this
is done before the DECODE2 phase, it will cause the
DISPATCH phase to jump to an address that is 1MB past the
10 Dispatch Table Entry that should have been used. This is
an illegal address, and will cause an exception when the
DISPATCH is attempted. The handler for this exception
will notice that the bit in disp_table had been set, and
will cause an Address Error exception to be generated.
15 This method has less overhead than the first method when
the pc is even, but significantly more overhead when the
pc is odd. The single POWER used by this method is as
follows.

```
rlimi disp_table,pc,20,0x00100000
```

20

MOVEM Register List Optimizations

The 68020 MOVEM instruction has a 16 bit register
list mask associated with it. There is one bit for each
of the 16 A/D registers, indicating if the register
25 should be moved. In general, the emulator would need to
perform 16 individual tests of the bits to determine if
the corresponding register needed to be read/written.
This can be very time consuming. Due to compiler
register allocation, parameter passing, and calling
30 conventions, there are some registers that rarely appear
in the register list of the MOVEM instruction, and others
that are much more frequent. Taking advantage of this,
the emulator can first check to see if any of the

infrequent registers are on the list, using a bitwise
 "and" operation against the mask. If none of those bits
 were set, then there is a much smaller list of bits that
 need to be checked individually (7 instead of 16), which
 5 will improve performance. The infrequent registers that
 the emulator currently test for are d0-d3/a0-a1/a5-a7 and
 the frequent registers are d4-d7/a2-a4.

Optimizations based on Opcode Synonyms

10 In the 68020 instruction set, there are many cases
 where two different opcode encodings perform exactly the
 same operation. Since the emulator can uniquely decode
 each of the 65536 possible opcode encodings, it can make
 sure that the dispatch table entries for two opcode
 15 synonyms are the same, and have exactly the same
 performance. Below is a list of instructions and their
 synonyms.

	add.<b,w,l>	#imm,dn	≈	addi.<b,w,l>	#imm,dn
	adda.w	#imm,an	≈	lea.l	imm(an),an
20	addq.w	#imm,an	≈	addq.l	#imm,an
	and.<b,w,l>	#imm,dn	≈	andi.<b,w,l>	#imm,dn
	asl.<b,w,l>	#1,dn	≈	add.<b,w,l>	dn,dn
	bra.s	* +4	≈	dbt.w	dx,xxxx (32 bit nop)
	bra.w	d16	≈	jmp	d16(pc)
25	bsr.w	d16	≈	jsr	d16(pc)
	clr.l	dn	≈	moveq.l	#0,dn
	cmp.<b,w,l>	#imm,dn	≈	cmpi.<b,w,l>	#imm,dn
	lea.l	(as),ad	≈	movea.l	as,ad
	lea.l	abs.w,an	≈	movea.w	#imm16,an
30	lea.l	abs.l,an	≈	movea.l	#imm32,an
	movea.l	(a7),a7	≈	unlk	a7
	or.<b,w,l>	#imm,dn	≈	ori.<b,w,l>	#imm,dn

sub.<b,w,l> #imm,dn ≈ subi.<b,w,l> #imm,dn
 subq.w #imm,an ≈ subq.l #imm,an

Optimizations based on Operands

5 In many cases, opcodes that have the same register specified as both the source and destination behave the same as some other opcode which executes faster. For these opcodes, we create an optimized dispatch table entry which is the same as the simpler opcode, instead of using the
 10 dispatch table entry for the general case. The table below shows the transformations.

	and.<b,w,l>	dn,dn	->	tst.<b,w,l>	dn
	cmp.<b,w,l>	dn,dn	->	tst.l	zero
	cmpa.l	an,an	->	tst.l	zero
15	eor.<b,w,l>	dn,dn	->	clr.<b,w,l>	dn
	exg.l	rn,rn	->	nop	
	lea.l	(an),an	->	nop	
	move.<b,w,l>	dn,dn	->	tst.<b,w,l>	dn
	movea.l	an,an	->	nop	
20	movea.<w,l>	(an)+,an	->	movea.<w,l>	(an),an
	or.<b,w,l>	dn,dn	->	tst.<b,w,l>	dn
	sub.<b,w,l>	dn,dn	->	clr.<b,w,l>	dn (also clear ccr.x)

25 In many cases, memory to memory MOVE instructions that have the same source and destination addresses behave the same as a TST instruction. If we can assume that the read and the write will be to the same RAM location, and that there are no side effects (write protected or I/O space
 30 accesses), then it should be safe to omit the write, and just do the read. The table below shows the transformations.

```

move.<b,w,l> (an),(an)    ->  tst.<b,w,l>    (an)
move.<b,w,l> -(an),(an)   ->  tst.<b,w,l>   -(an)
move.<b,w,l> (an)+,-(an)  ->  tst.<b,w,l>    (an)
move.<b,w,l> -(an),(an)+  ->  tst.<b,w,l>   <1,2,4> (an)
    
```

5

Optimizations based on repeated Opcodes

There are many places in the Macintosh QuickDraw routines where data movement loops are "unwound" and contain 16 repeated sequences of the same instruction, followed by a DBRA looping instruction. There are also cases of shorter sequences in compiler generated code for structure copying. Since the emulation of a repeated sequence of 68020 opcodes will cause a repeated sequence of POWER instructions to be executed within the emulator, it is possible for the emulation of one of these opcodes to detect that the next opcode is the same, and eliminate some of the decoding and dispatching overhead, which will improve the performance of the subsequent instances of the same opcode. If instruction tracing or an interrupt is pending, this optimization cannot be performed, because special event processing would need to occur at the instruction boundaries. The opcodes that the emulator currently detects repeated sequence of are shown below.

```

                move.l    d6,(a1)+
25             move.l    d6,(a5)+
                eor.l     d6,(a5)+
                move.l    (a0)+,(a1)+
                move.l    (a0)+,(a2)+
                move.l    (a1)+,(a0)+
30             move.l    (a2)+,(a1)+
                move.l    (a4)+,(a5)+
    
```

Optimizations based on Opcode Sequences

In compiler generated code, and sometimes in the Macintosh ROM code there are common sequences of two and sometimes three instructions that occur frequently due to runtime calling conventions. In many cases detecting these sequences is done by having the emulation of the first instruction of the sequence check to see if it is followed by the second instruction of the sequence. For sequences of three instructions, if the pair of the first two had been detected, then the check for the third instruction can be made. The emulator currently detects and optimizes the following pairs and triples of instructions. As with other optimizations of this kind, the optimization cannot be performed if special events are pending. The sequences that the emulator may detect are shown below.

```

5      move.b (a1)+,(a0)+          bne.s  *-2
      move.b (a0)+,d0             cmp.b (a1)+,d0
      move.b (a0)+,d1             cmp.b (a1)+,d1
      move.l (a7)+,(a7)           rts
20     move.l ([d16,ZA0,ZA0.w*1],d16),-(a7) . rts
      move.l abs.w,-(a7)          rts
      movem.l d16(a6),reg_list    unlk   a6          rts
      movem.l (a7)+,reg_list     unlk   a6          rts

```

25 ATrap Dispatcher Acceleration

The Macintosh OS and Toolbox use the unimplemented LineA opcode space to encode system calls. The LineA dispatcher in the ROM executes several instructions (on every call) to dispatch to the desired routine. This whole process can be greatly improved by having the emulator directly dispatch these instructions. However, there must also be a way for this feature to be disabled if the standard LineA dispatcher has been replaced, such as when

using the A-Trap record/break features of MacsBug. In order to achieve this compatibility, we need to know if the handler address (in vector offset \$28) has changed. Fortunately, we do not have to worry about tracing, since
5 the 68000 will not take a trace trap on a LineA (or any other unimplemented) instruction.

IV. Conclusion

This emulation technique is also optimized for an
10 upgraded version of the POWER architecture, referred to as the POWER PC architecture. This updated version of the POWER architecture executes the same opcodes listed above for the POWER architecture instructions, except with different mnemonics.

15 Thus, a highly efficient method for decoding 68020 instructions in the POWER architecture processor, utilizing a special table of instruction sequences, specific sequences of instructions and other optimizations is provided. The alignment and method of indexing the table
20 contribute to the high speed decoding of the 68020 instructions. In particular, a single POWER instruction can be used to index the table, by aligning the dispatch table on a 512K byte boundary within a two megabyte range, shifting the opcode left by three bits and inserting the
25 shifted opcode into the address of the base of the table. Each entry in the table contains two POWER instructions. The first instruction generally performs an operation specific to source addressing mode, and the second instruction branches to an emulation routine which is
30 generally independent of the source addressing mode. This allows reduced code size, and better memory reference and cache locality for the emulation routines.

There are four phases of the instruction decoding process. DECODE1 forms the address of the next dispatch

table entry by shifting and inserting the next 68020 opcode into the address of the base of the dispatch table. DECODE2 moves that dispatch table address to the power architecture CTR register. PREFETCH reads the 16 bits
5 (assigned extended to 32 bits) that follow the next 68020 instruction into a temporary holding register called `prefetch_data`. Finally, DISPATCH jumps to the address that was computed by DECODE1.

This sequence of four instruction macros is the
10 minimum number of instructions in this embodiment of the invention that this task can be performed in on the POWER architecture, and the use of opcode prefetching reduces or eliminates pipeline stalls due to memory reference latency.

Overall, performance of an emulating guest
15 instructions is enhanced by reducing the number of host instructions needed to be executed to emulate the guest instruction. One primary reason for this benefit is the direct access to the instructions in the dispatch table, combined with prefetching of the guest instructions.
20 Furthermore, opcode sequences of any length may be optimized, and decoding and dispatch logic bypassed for the sequence, to further enhance emulator performance.

The foregoing description of preferred embodiments of the present invention has been provided for the purposes of
25 illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Obviously, many modifications and variations will be apparent to practitioners skilled in this art. The embodiments were chosen and described in order to best
30 explain the principles of the invention and its practical application, thereby enabling others skilled in the art to understand the invention for various embodiments and with various modifications as are suited to the particular use contemplated. It is intended that the scope of the

invention be defined by the following claims and their equivalents.

CLAIMS

What is claimed is:

- 1 1. For a host processor which executes host
2 instructions and includes host addressable memory, a system
3 for decoding guest instructions, comprising:
4 a guest program store in host addressable memory to
5 store at least one program of guest instructions;
6 an emulation program store in the host addressable
7 memory having a set of emulation programs, at least a
8 subset of the set of emulation programs including
9 respective host instruction routines for respective
10 particular guest instructions;
11 dispatch logic coupled with the guest program store
12 and the emulation program store to dispatch emulation
13 programs in the set, in response to guest instructions in
14 the at least one program of guest instructions and in
15 response to the emulation programs in the set of emulation
16 programs; and
17 emulation logic, comprising host instructions embedded
18 within a particular emulation program in the set of
19 emulation programs, to detect a particular sequence of
20 guest instructions, and in response to detection of the
21 particular sequence to bypass the dispatch logic for guest
22 instructions in the particular sequence.
- 1 2. The system of claim 1, wherein the particular
2 sequence comprises repeated guest instructions.
- 1 3. The system of claim 1, wherein the particular
2 sequence comprises more than two guest instructions.

1 4. The system of claim 1, wherein the dispatch logic
2 comprises host instructions embedded within emulation
3 programs in the set of emulation programs.

1 5. The system of claim 1, wherein the emulation
2 programs store comprises:

3 a dispatch store in host processor addressable memory
4 having a set of dispatch entries, each dispatch entry in
5 the set including a plurality of host instructions of the
6 emulation program corresponding to a guest instruction;

7 an emulation routine store in host addressable memory
8 having a set of emulation entries beginning at
9 corresponding emulation table addresses, each emulation
10 entry in the set including a host instruction routine for
11 the emulation program;

12 wherein the plurality of host instructions in a subset
13 of the set of dispatch entries includes a host jump
14 instruction which causes a jump upon execution by the host
15 processor to an emulation routine table address of a
16 corresponding emulation entry in the emulation routine
17 table; and

18 wherein the host instruction routines in a subset of
19 the set of emulation entries include host instructions
20 which upon execution by the host processor form an
21 emulation program address to a dispatch entry in response
22 to a next guest instruction and jump directly to the
23 dispatch table address;

24 such that emulation programs in the emulation program
25 store comprise a plurality of instructions in the dispatch
26 table and a host instruction routine in the emulation
27 routine store.

1 6. The system of claim 5, further including:
2 a guest instruction pointer store for a guest
3 instruction pointer indicating a guest instruction address
4 in the program of guest instructions;
5 a prefetched guest instruction store for a guest
6 instruction read from the program of guest instructions in
7 response to the guest instruction pointer; and
8 an emulation program pointer store for an emulation
9 program address formed in response to the current guest
10 instruction read from the prefetched guest instruction
11 store.

1 7. The system of claim 6, wherein emulation programs
2 in a subset of the set of emulation programs include:
3 a first segment of host instructions which upon
4 execution by the host processor forms an emulation program
5 address in the emulation program pointer store in response
6 to the current guest instruction in the prefetch guest
7 instruction store;
8 a second segment of host instructions which upon
9 execution by the host processor reads a next guest
10 instruction from an address indicated by the guest
11 instruction pointer into the prefetch guest instruction
12 store;
13 a third segment of host instructions which upon
14 execution by the host processor causes a jump to the
15 emulation program indicated by the emulation program
16 address in the emulation program pointer store; and
17 a fourth segment of host instructions which upon
18 execution by the host processor bypass at least the third
19 segment upon detection of the particular sequence for guest
20 instructions within the particular sequence.

1 8. For a host processor which executes host
2 instructions and includes host addressable memory, a system
3 for decoding guest instructions, comprising:
4 a guest program store in host addressable memory to
5 store at least one program of guest instructions;
6 an emulation program store in the host addressable
7 memory having a set of emulation programs, at least a
8 subset of the set of emulation programs including
9 respective host instruction routines for respective
10 particular guest instructions;
11 prefetch logic comprising host instructions embedded
12 within emulation programs in the set of emulation programs,
13 to prefetch guest instructions in the at least one program
14 of guest instructions to a prefetch store;
15 dispatch logic comprising host instructions embedded
16 within emulation programs in the set of emulation programs,
17 to dispatch emulation programs in the set, in response to
18 guest instructions in the prefetch store; and
19 emulation logic, comprising host instructions embedded
20 within a particular emulation program in the set of
21 emulation programs, responsive to guest instructions in the
22 prefetch store to detect a particular sequence of guest
23 instructions, and in response to detection of the
24 particular sequence to bypass the dispatch logic for guest
25 instructions in the particular sequence.

1 9. The system of claim 8, wherein the particular
2 sequence comprises repeated guest instructions.

1 10. The system of claim 8, wherein the particular
2 sequence comprises more than two guest instructions.

1 11. The system of claim 8, wherein the emulation
2 programs store comprises:
3 a dispatch store in host processor addressable memory
4 having a set of dispatch entries, each dispatch entry in
5 the set including a plurality of host instructions of the
6 emulation program corresponding to a guest instruction;
7 an emulation routine store in host addressable memory
8 having a set of emulation entries beginning at
9 corresponding emulation table addresses, each emulation
10 entry in the set including a host instruction routine for
11 the emulation program;
12 wherein the plurality of host instructions in a subset
13 of the set of dispatch entries includes a host jump
14 instruction which causes a jump upon execution by the host
15 processor to an emulation routine table address of a
16 corresponding emulation entry in the emulation routine
17 table; and
18 wherein the host instruction routines in a subset of
19 the set of emulation entries include host instructions
20 which upon execution by the host processor form an
21 emulation program address to a dispatch entry in response
22 to a next guest instruction and jump directly to the
23 dispatch table address;
24 such that emulation programs in the emulation program
25 store comprise a plurality of instructions in the dispatch
26 table and a host instruction routine in the emulation
27 routine store.

1 12. The system of claim 11, further including:
2 a guest instruction pointer store for a guest
3 instruction pointer indicating a guest instruction address
4 in the program of guest instructions; and
5 an emulation program pointer store for an emulation
6 program address formed in response to the current guest
7 instruction read from the prefetched guest instruction
8 store.

1 13. The system of claim 12, wherein emulation
2 programs in a subset of the set of emulation programs
3 include:
4 a first segment of host instructions which upon
5 execution by the host processor forms an emulation program
6 address in the emulation program pointer store in response
7 to the guest instruction in the prefetch store;
8 a second segment of host instructions which upon
9 execution by the host processor reads a next guest
10 instruction from an address indicated by the guest
11 instruction pointer into the prefetch store;
12 a third segment of host instructions which upon
13 execution by the host processor causes a jump to the
14 emulation program indicated by the emulation program
15 address in the emulation program pointer store; and
16 a fourth segment of host instructions which upon
17 execution by the host processor bypass at least the third
18 segment upon detection of the particular sequence for guest
19 instructions within the particular sequence.

1 14. For a host processor which executes host
2 instructions and includes host addressable memory, a system
3 for decoding guest instructions, comprising:
4 a guest program store in host addressable memory to
5 store at least one program of guest instructions;
6 a dispatch store in host processor addressable memory
7 having a set of dispatch entries, each dispatch entry in
8 the set including a plurality of host instructions of an
9 emulation program corresponding to a guest instruction;
10 an emulation routine store in host addressable memory
11 having a set of emulation entries beginning at
12 corresponding emulation table addresses, each emulation
13 entry in the set including a host instruction routine for
14 an emulation program corresponding to a guest instruction;
15 prefetch logic comprising host instructions embedded
16 within emulation programs, to prefetch guest instructions
17 in the at least one program of guest instructions to a
18 prefetch store;
19 dispatch logic comprising host instructions embedded
20 within emulation programs, to dispatch emulation programs
21 in response to guest instructions in the prefetch store;
22 and
23 emulation logic, comprising host instructions embedded
24 within a particular emulation program, responsive to guest
25 instructions in the prefetch store to detect a particular
26 sequence of guest instructions, and in response to
27 detection of the particular sequence to bypass the dispatch
28 logic for guest instructions in the particular sequence.

1 15. The system of claim 14, wherein the particular
2 sequence comprises repeated guest instructions.

1 16. The system of claim 14, wherein the particular
2 sequence comprises more than two guest instructions.

1 17. The system of claim 14,
2 wherein the plurality of host instructions in a subset
3 of the set of dispatch entries includes a host jump
4 instruction which causes a jump upon execution by the host
5 processor to an emulation routine table address of a
6 corresponding emulation entry in the emulation routine
7 table; and
8 wherein the dispatch logic in a subset of the set of
9 emulation entries includes host instructions which upon
10 execution by the host processor form an emulation program
11 address to a dispatch entry in response to a next guest
12 instruction in the prefetch store and jump directly to the
13 dispatch table address.

1 18. The system of claim 17, further including:
2 a guest instruction pointer store for a guest
3 instruction pointer indicating a guest instruction address
4 in the program of guest instructions; and
5 an emulation program pointer store for an emulation
6 program address formed in response to the current guest
7 instruction read from the prefetched guest instruction
8 store.

1 19. The system of claim 18, wherein emulation
2 programs in a subset of the set of emulation programs
3 include:
4 a first segment of host instructions which upon
5 execution by the host processor forms an emulation program
6 address in the emulation program pointer store in response
7 to the guest instruction in the prefetch store;

8 a second segment of host instructions which upon
9 execution by the host processor reads a next guest
10 instruction from an address indicated by the guest
11 instruction pointer into the prefetch store;

12 a third segment of host instructions which upon
13 execution by the host processor causes a jump to the
14 emulation program indicated by the emulation program
15 address in the emulation program pointer store; and

16 a fourth segment of host instructions which upon
17 execution by the host processor bypass at least the third
18 segment upon detection of the particular sequence for guest
19 instructions within the particular sequence.

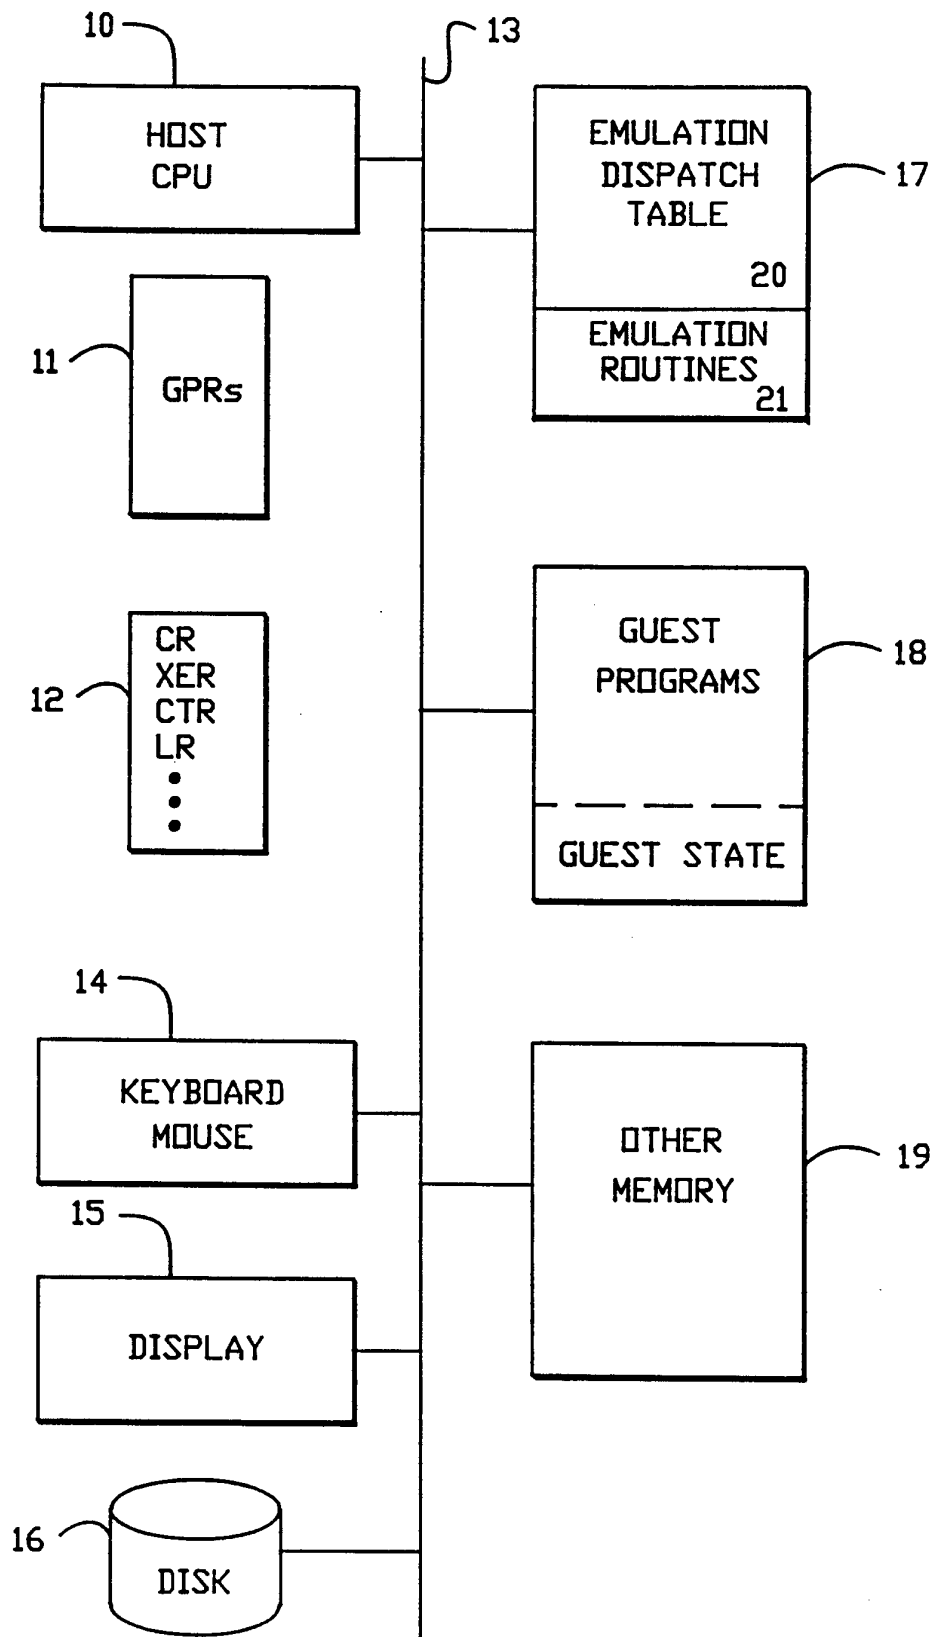


FIG.-1
1/9

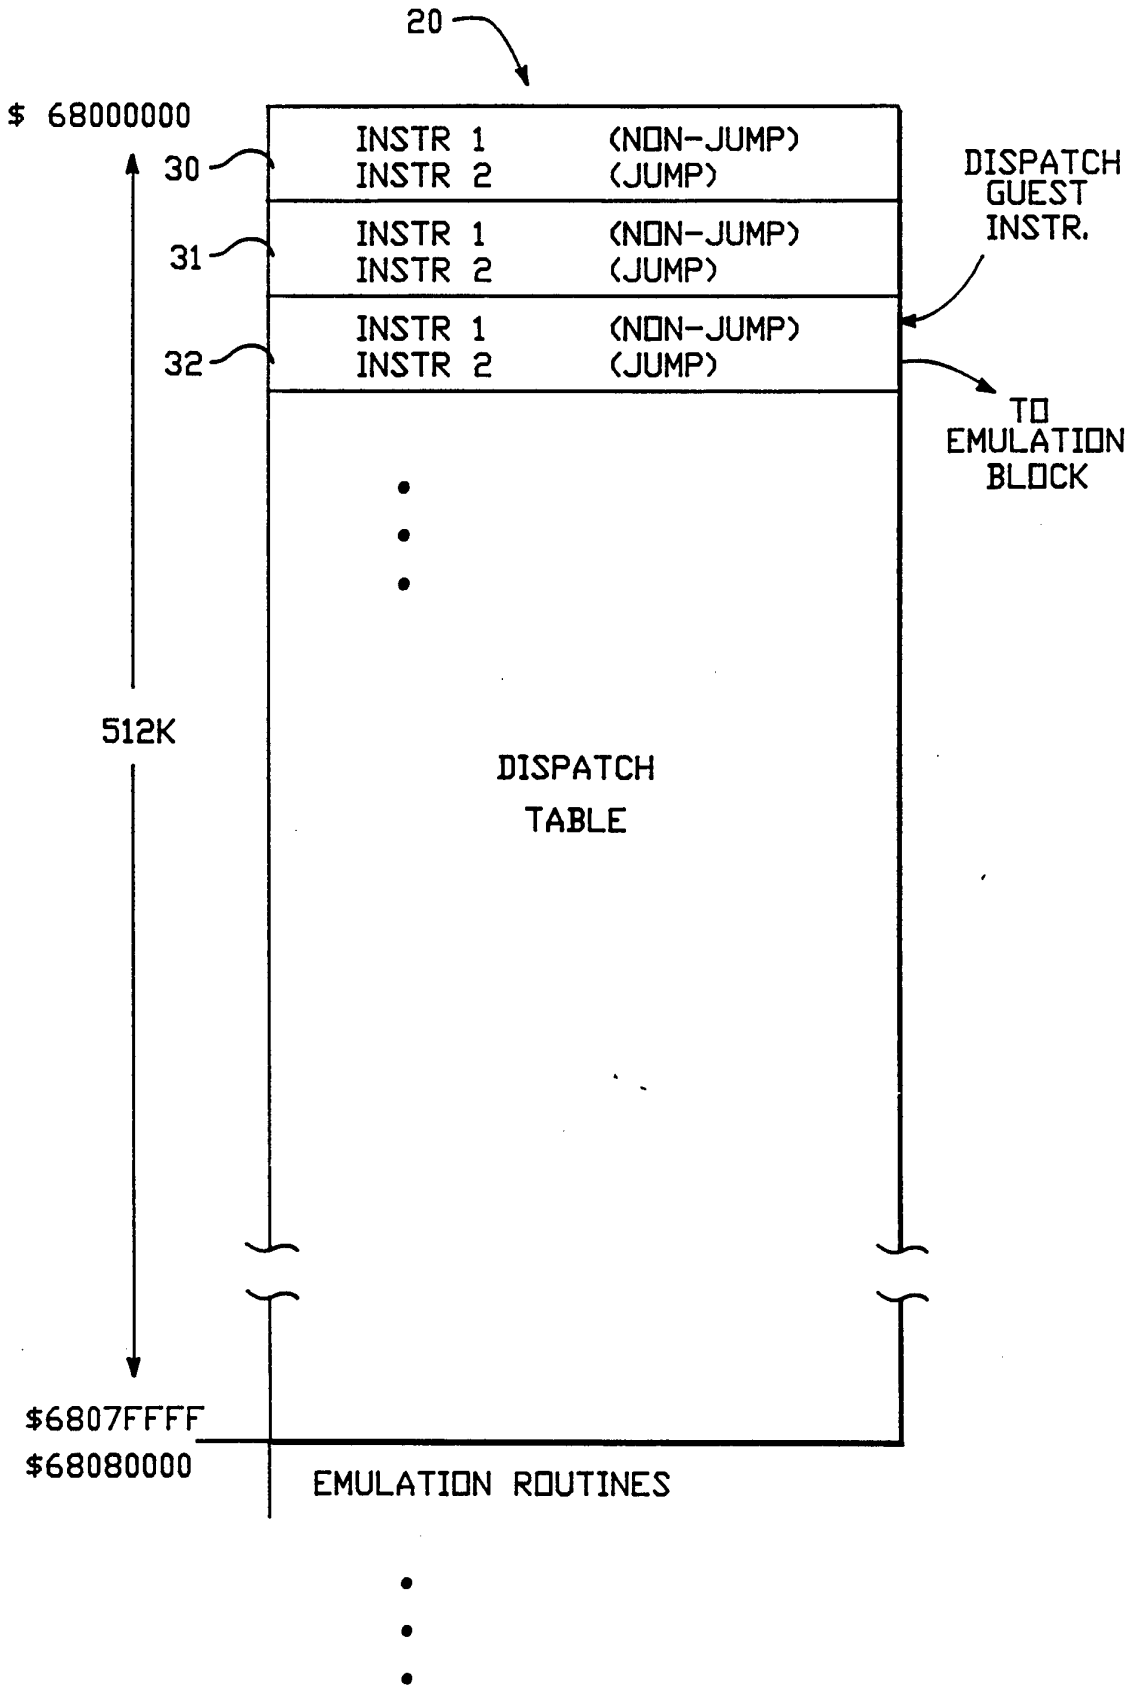
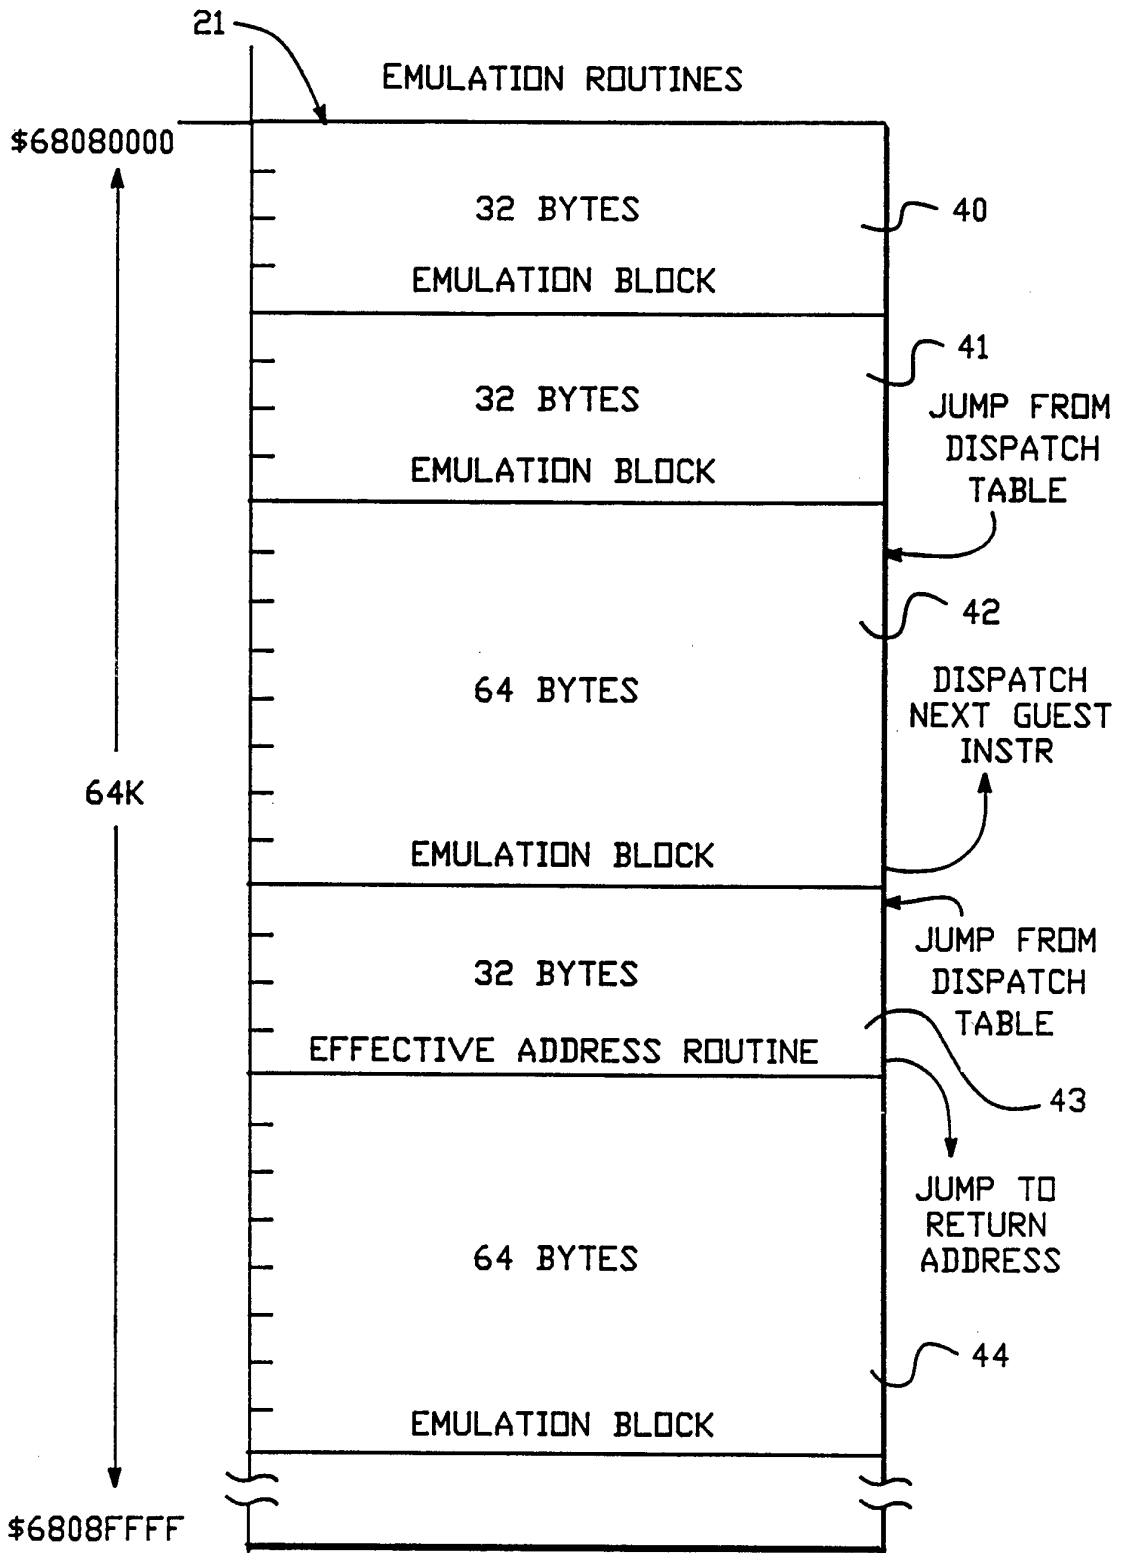


FIG.-2



21

FIG.-3

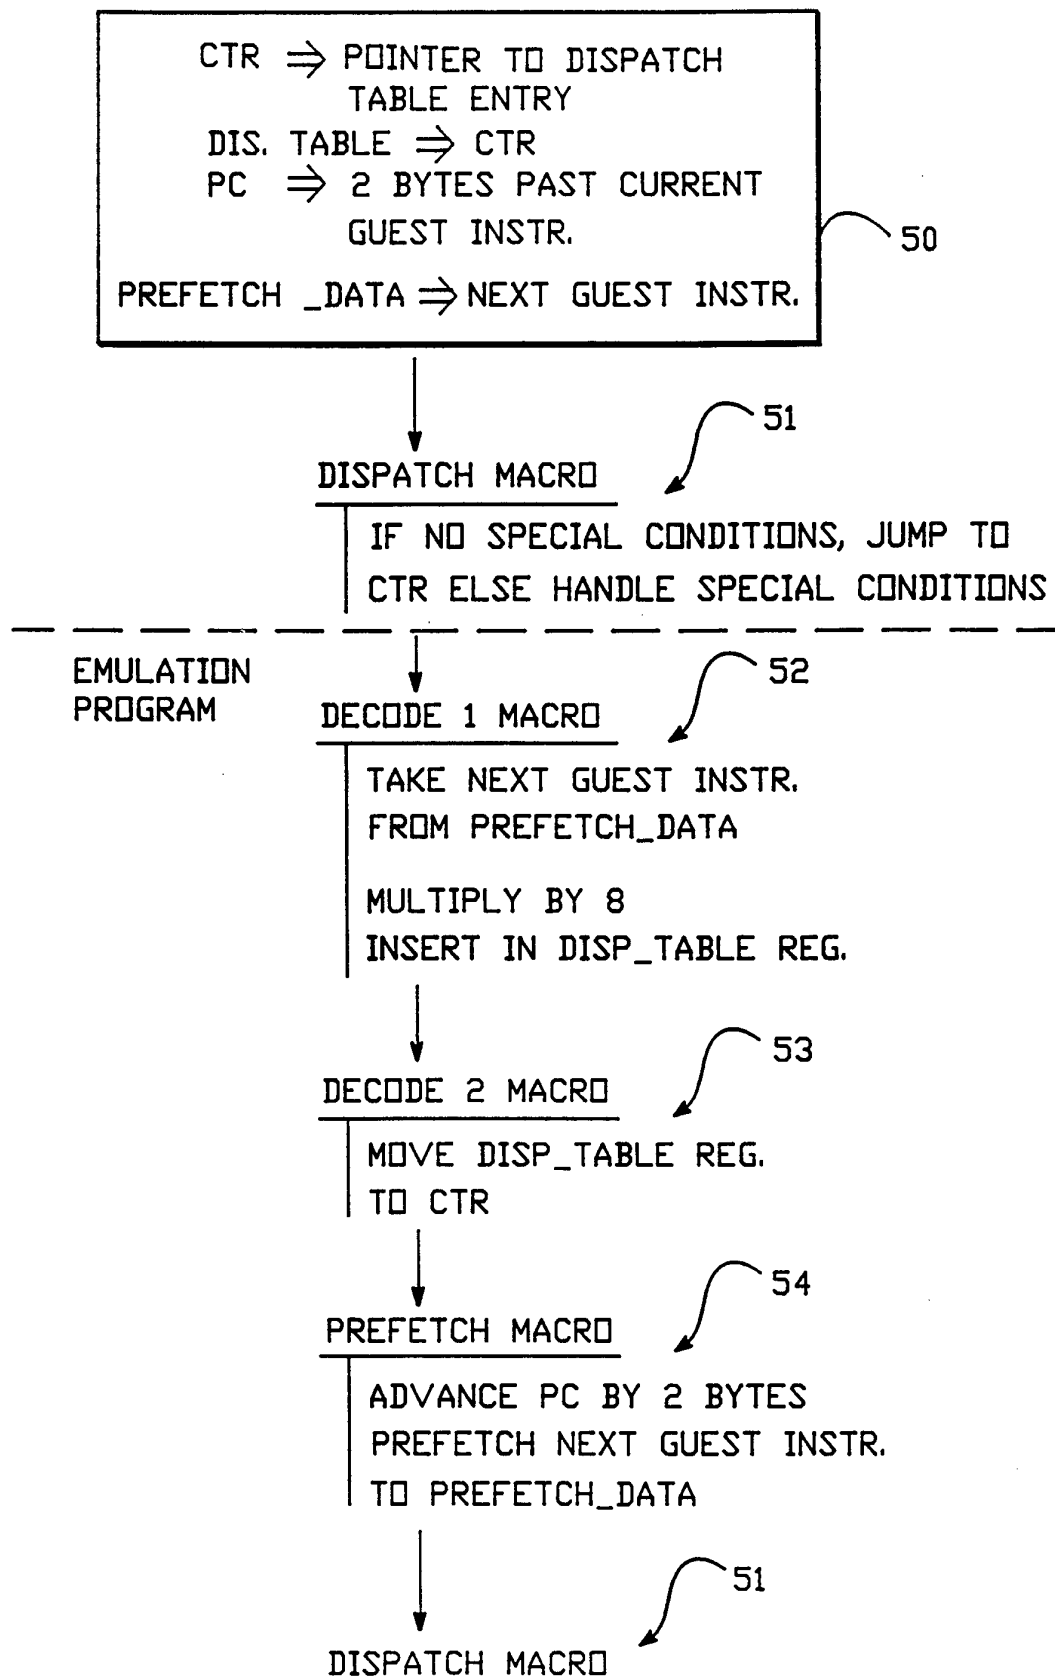


FIG. -4

1	INST 1 (NON-JUMP-ADDRESSING MODE OP)	
2	INST 2 (JUMP-TO EMULATION BLOCK)	DISPATCH TABLE
<hr/>		
3	INST(S) (START EMULATION BLOCK)	
4	DECODE1 MACRO	
5	INST(S) (EMULATION BLOCK)	EMULATION CODE TABLE
6	DECODE2 MACRO	
7	INST(S) (EMULATION BLOCK)	
8	PRE-FETCH MACRO	
9	INST(S) (EMULATION BLOCK)	
10	DISPATCH MACRO	

FIG.-5

1	INST 1 (NON-JUMP-SET RETURN ADDR)	
2	INST 2 (JUMP-TO EFFECTIVE ADDR ROUTINE)	DISPATCH TABLE
<hr/>		
3	INST(S) (EFFECTIVE ADDR ROUTINE)	
4	INST (MOVE rtn_addr TO lr)	
5	INST(S) (JUMP TO lr)	EMULATION CODE TABLE
6	INST(S) (START EMULATION BLOCK)	
7	DECODE1 MACRO	
8	INST(S) (EMULATION BLOCK)	
9	DECODE2 MACRO	
10	PRE-FETCH MACRO	
11	INST(S) (EMULATION BLOCK)	
12	DISPATCH MACRO	

FIG.-6

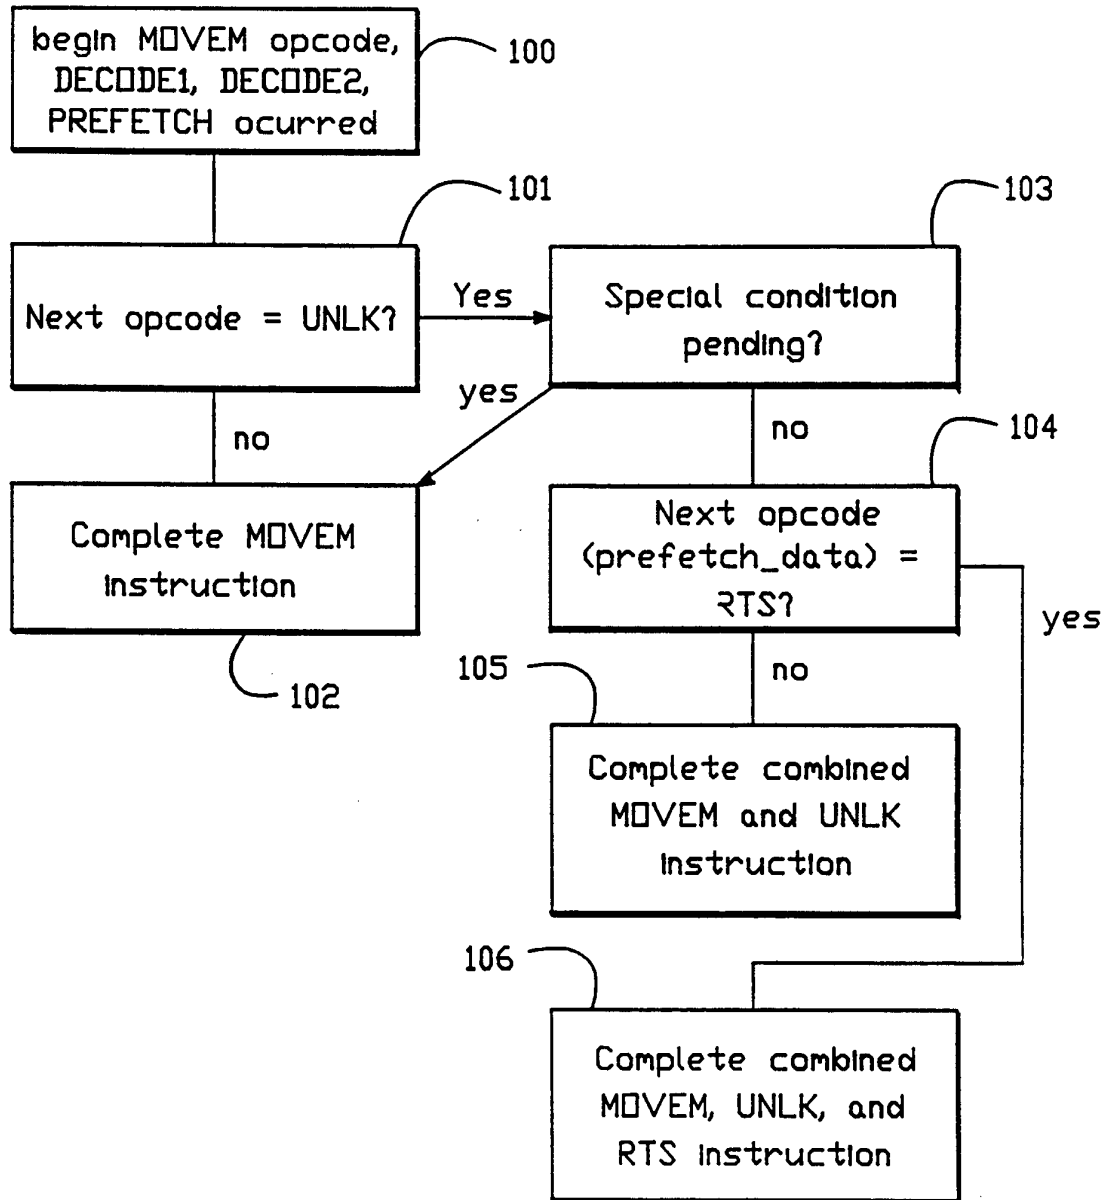


FIG.-7

1st GUEST	1	INST 1
	2	INST 2
	3	INST(S) (OPCODE EXTENSION NEEDED)
2nd GUEST	4	PREFETCH FROM PC TO GPR (ADV PC)
	5	TEST GPR FOR SEQUENCE (SAVE)
	6	USE EXTENSION IN PREFETCH_DATA
	7	INST(S)
	8	DECODE1 MACRO ON GPR
	9	DECODE2 MACRO ON GPR
	10	INST(S)
3rd GUEST	11	PREFETCH MACRO INTO PREFETCH DATA
	12	CHECK FOR SPECIAL CONDITIONS
	13	BRANCH IF SEQUENCE AND NO SPECIAL CONDITION TO 16
	14	INST(S)
	15	DISPATCH MACRO (FOR 2nd GUEST)
	16	INST(S)
	17	TEST PREFETCH_DATA FOR SEQUENCE
	18	BRANCH IF SEQUENCE TO 25
	19	INST(S)
	20	DECODE1 MACRO
	21	DECODE2 MACRO
	22	INST(S)
4th GUEST	23	PREFETCH MACRO
	24	DISPATCH MACRO (FOR 3rd GUEST)
	25	INST(S)
4th GUEST	26	FETCH OPCODE FOR RTS
	27	DECODE1 MACRO
	28	DECODE2 MACRO
	29	INST(S)
5th GUEST	30	PREFETCH MACRO
	31	PREFETCH MACRO (FOR 4th GUEST RETURN)

FIG.—8

7/9

SUBSTITUTE SHEET (RULE 26)

1st GUEST	1	INST 1
	2	INST 2
<hr/>		
2nd GUEST	3	TEST PREFETCH_DATA FOR SEQUENCE
	4	BRANCH IF SEQUENCE AND NO SPECIAL COND. TO 11
	5	INST(S)
	6	DECODE1 MACRO
	7	DECODE2 MACRO
	8	INST(S)
3rd GUEST	9	PREFETCH MACRO
	10	DISPATCH MACRO (2nd GUEST)
3rd GUEST	11	PREFETCH MACRO
	12	INST(S)
	13	DECODE1 MACRO
	14	INST(S)
	15	DECODE2 MACRO
4th GUEST	16	PREFETCH MACRO
	17	DISPATCH MACRO (3rd GUEST)

FIG. - 9

1st GUEST	1	TEST FOR REPEAT
	2	JUMP TO 3
<hr/>		
	3	INST(S)
	4	DECODE1 MACRO
(3rd GUEST)	5	PREFETCH MACRO
	6	INST(S)
	7	BRANCH IF REPEAT TO 11
	8	DECODE2 MACRO
	9	INST(S)
	10	DISPATCH MACRO ((2+N)th GUEST)
	11	TEST FOR REPEAT
	12	BRANCH IF SPECIAL CONDITION TO 8
	13	INST(S)
	14	DECODE1 MACRO
((3+N)th GUEST)	15	PREFETCH MACRO
	16	INST(S)
	17	BRANCH IF REPEAT TO 11
	18	BRANCH TO 8

FIG.—10

INTERNATIONAL SEARCH REPORT

 Intern. Patent Application No
 PCT/US 94/03862

 A. CLASSIFICATION OF SUBJECT MATTER
 IPC 5 G06F9/318 G06F9/455

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

B. FIELDS SEARCHED

 Minimum documentation searched (classification system followed by classification symbols)
 IPC 5 G06F

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

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EDN ELECTRICAL DESIGN NEWS vol. 20, no. 6, March 1975, NEWTON, MASSACHUSETTS US pages 61 - 66 R. E. VAHLSTROM 'Microprogramming Emulation of a General-Purpose Processor' see page 63, right column, line 30 - page 64, left column, line 5 ---	1, 8, 14
A	US, A, 3 698 007 (MALCOLM ET AL.) 10 October 1972 see the whole document --- -/--	1, 8, 14

 Further documents are listed in the continuation of box C.

 Patent family members are listed in annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

4 August 1994

Date of mailing of the international search report

16. 08. 94

Name and mailing address of the ISA

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

Authorized officer

Daskalakis, T

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 94/03862

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>COMPUTER DESIGN. vol. 25, no. 12 , June 1986 , LITTLETON, MASSACHUSETTS US pages 87 - 90 B. SHERMAN 'Attention to Basics Reduces Risc in Ada Compiler Choice' see page 89, right column, paragraph 2 -paragraph 4 -----</p>	1,8,14

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 94/03862

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US-A-3698007	10-10-72	NONE	
