



(11) (21) (C) **2,039,640**
(22) 1991/04/03
(43) 1991/11/05
(45) 2000/01/11

(72) Vassiliadis, Stamatis, US

(72) Blaner, Bartholomew, US

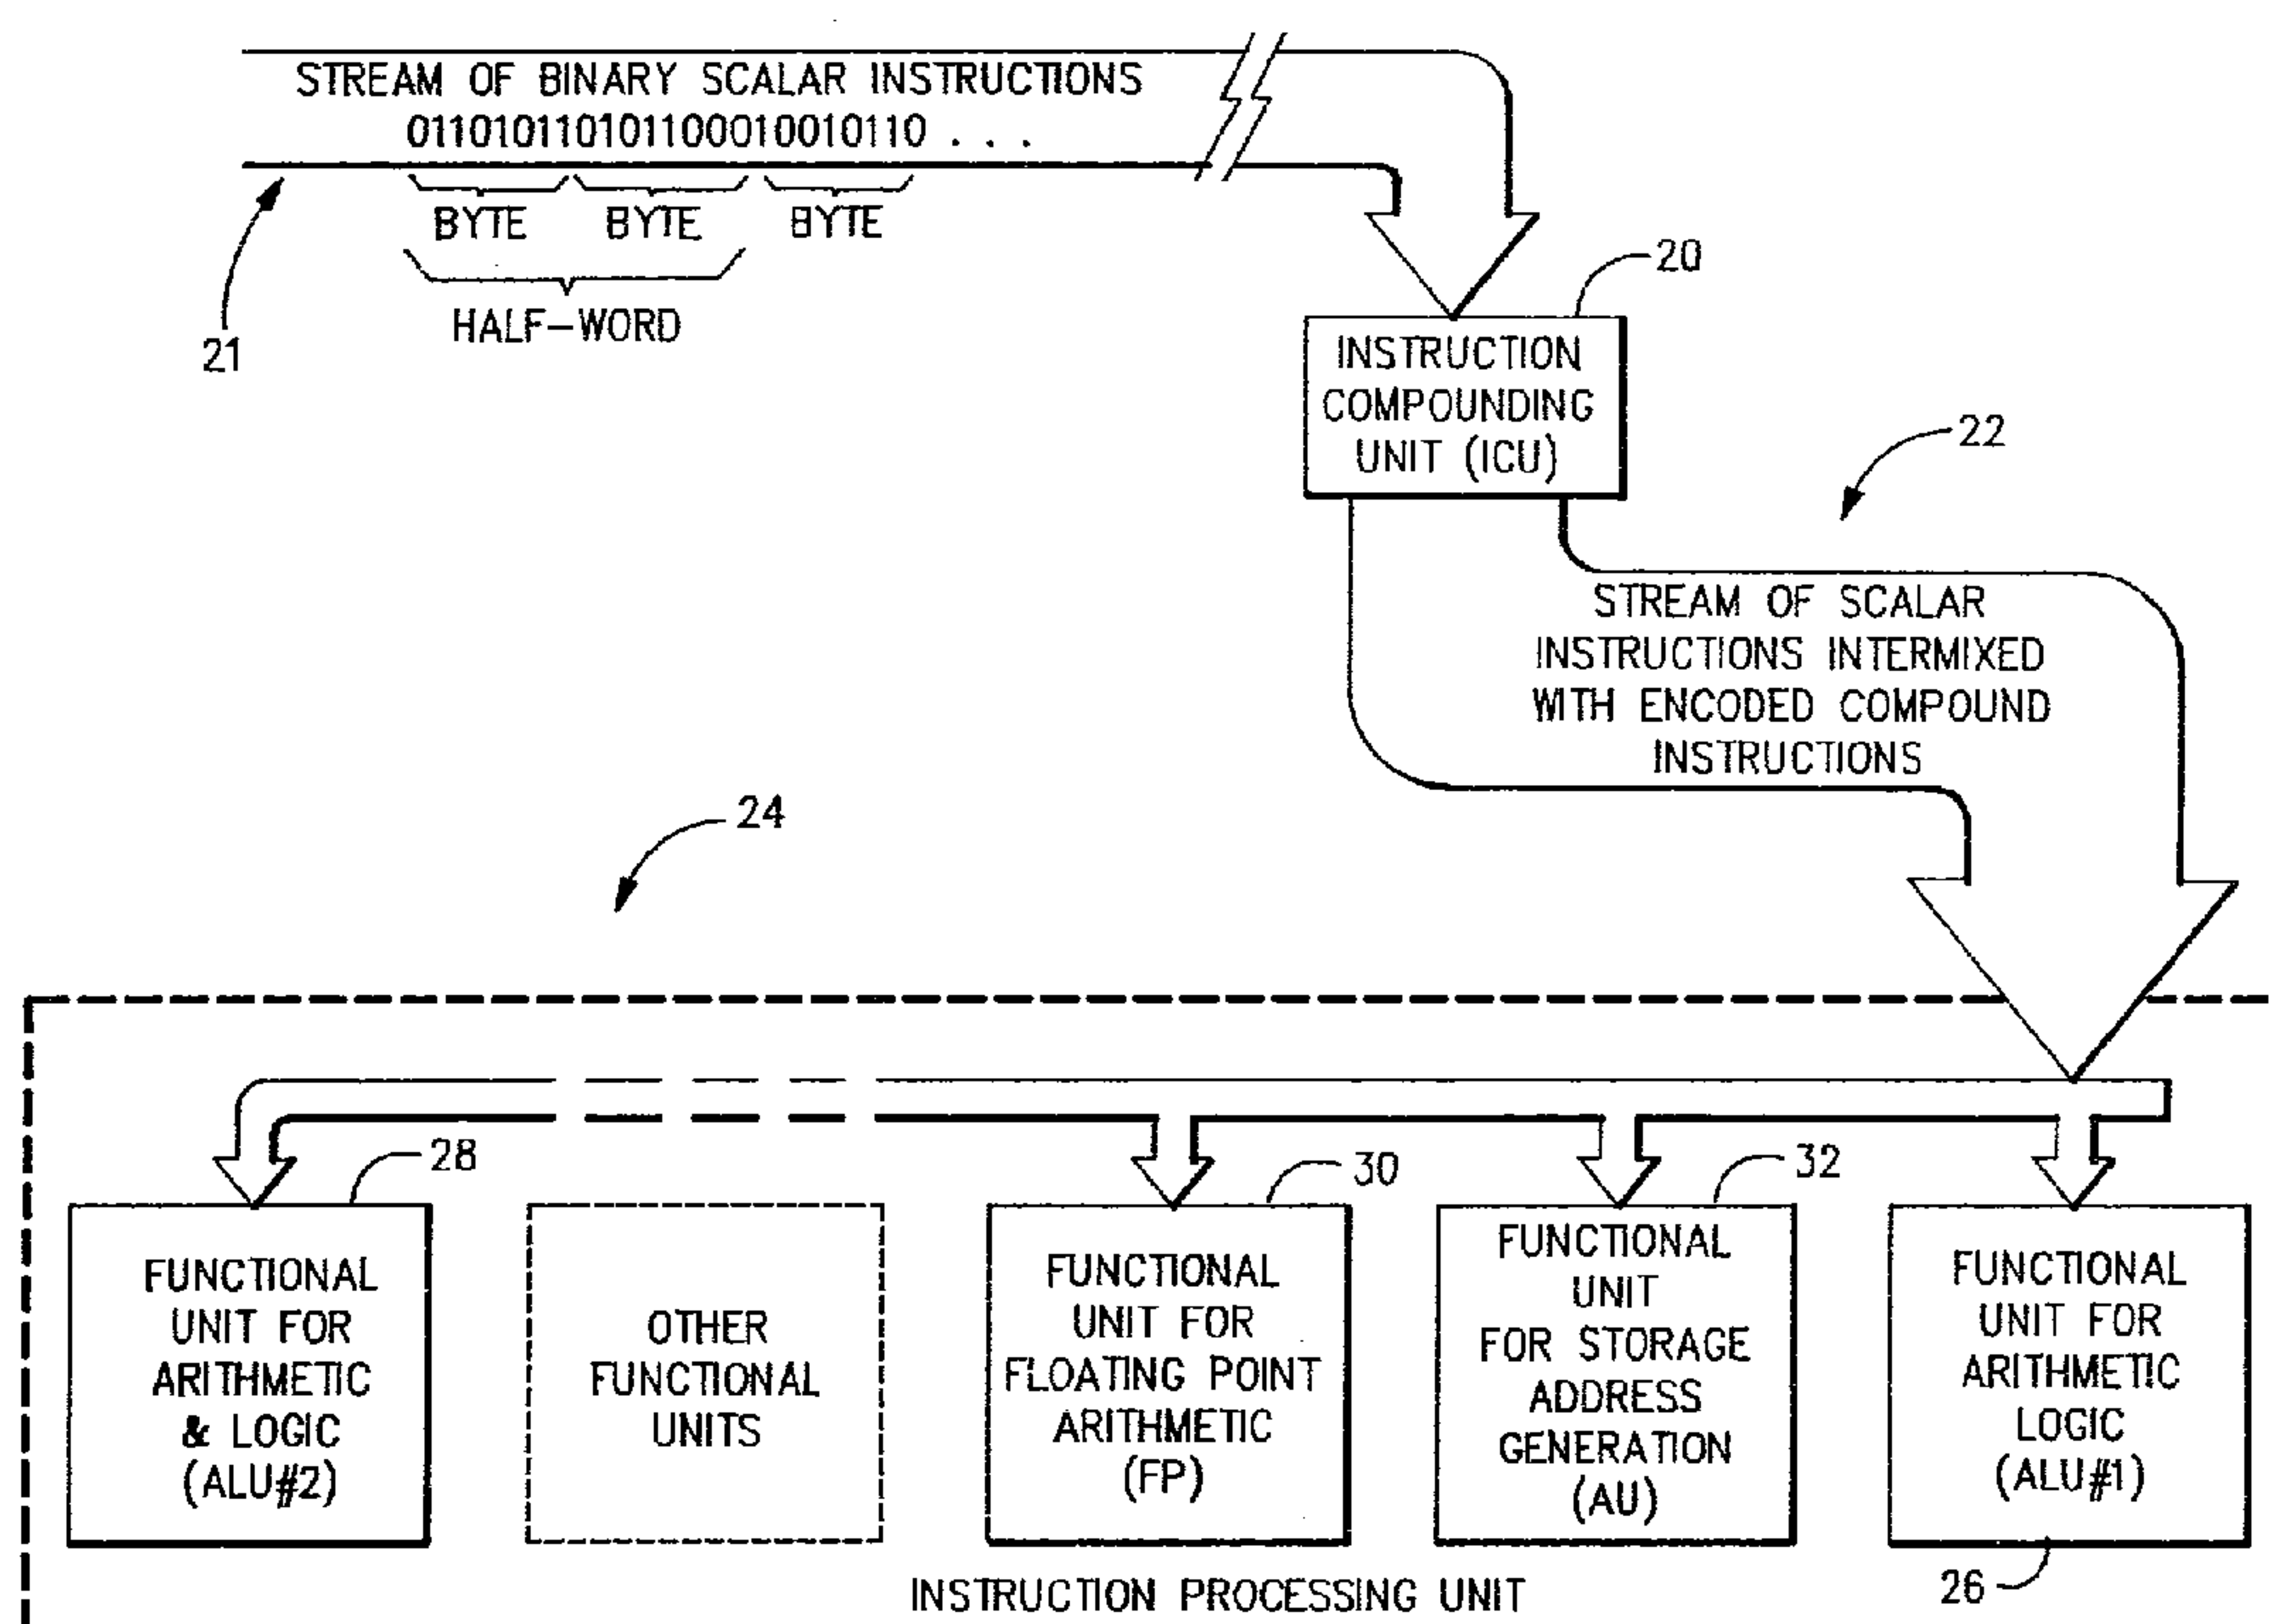
(73) INTERNATIONAL BUSINESS MACHINES CORPORATION, US

(51) Int.Cl.⁵ G06F 9/38

(30) 1990/05/04 (519,384) US

(54) **ARCHITECTURE EVOLUTIVE POUR MACHINE A JEU
D'INSTRUCTIONS COMPOSEES**

(54) **SCALABLE COMPOUND INSTRUCTION SET MACHINE
ARCHITECTURE**



(57) Scalable compound instruction set machine and method which provides for processing a set of instructions or program to be executed by a computer to determine statically which instructions may be combined into compound instructions which are executed in parallel by a scalar machine. Such processing looks for classes of instructions that can be executed in parallel without data-dependent or hardware-dependent interlocks. Without regard to their original sequence the individual instructions are combined with one or more other individual instructions to form a compound instruction which eliminates interlocks. Control information is appended to identify information relevant to the execution of the compound instructions. The result is a stream of scalar instructions compounded or grouped together before instruction decode time so that they are already flagged and identified for selective simultaneous parallel execution by execution units. The compounding does not change the object code results and existing programs realize performance improvements while maintaining compatibility with previously implemented systems for which the original set of instructions was provided.



SCALABLE COMPOUND INSTRUCTION SET MACHINE ARCHITECTURE

ABSTRACT

Scalable compound instruction set machine and method which provides for processing a set of instructions or program to be executed by a computer to determine statically which instructions may be combined into compound instructions which are executed in parallel by a scalar machine. Such processing looks for classes of instructions that can be executed in parallel without data-dependent or hardware-dependent interlocks. Without regard to their original sequence the individual instructions are combined with one or more other individual instructions to form a compound instruction which eliminates interlocks. Control information is appended to identify information relevant to the execution of the compound instructions. The result is a stream of scalar instructions compounded or grouped together before instruction decode time so that they are already flagged and identified for selective simultaneous parallel execution by execution units. The compounding does not change the object code results and existing programs realize performance improvements while maintaining compatibility with previously implemented systems for which the original set of instructions was provided.

SCALABLE COMPOUND INSTRUCTION SET
MACHINE ARCHITECTURE

This invention relates generally to parallel processing by computer, and more particularly relates to processing an instruction stream to identify those instructions which can be issued and executed in parallel in a specific computer system configuration.

10 The concept of parallel execution of instructions has helped to increase the performance of computer systems. Parallel execution is based on having separate functional units which can execute two or more of the same or different instructions simultaneously.

20 Another technique used to increase the performance of computer systems is pipelining. In general, pipelining is achieved by partitioning a function to be performed by a computer into independent subfunctions and allocating a separate piece of hardware, or stage, to perform each subfunction. Each stage is defined to occupy one basic machine cycle in time. Pipelining does provide a form of parallel processing since it is possible to execute multiple instructions concurrently. Ideally, one new instruction can be fed into the pipeline per cycle, with each instruction in the pipeline being in a different stage of execution. The operation is analogous to a manufacturing assembly line, with a number of instances of the manufactured product in varying stages of completion.

30 However, many times the benefits of parallel execution and/or pipelining are not achieved because of delays like those caused by data dependent interlocks and hardware dependent interlocks. An example of a data dependent interlock is a so-called write-read interlock where a first instruction must write its result before the second instruction can read and subsequently use it. An example of hardware dependent interlock is where a first instruction must use a particular hardware component and a second instruction must also use the same particular hardware component.

One of the techniques previously employed to avoid interlocks (sometimes called pipeline hazards) is called dynamic scheduling. Dynamic scheduling is based on the fact that with the inclusion of specialized hardware, it is possible to reorder instruction sequences after they have been issued into the pipeline for execution.

10 There have also been some attempts to improve performance through so-called static scheduling which is done before the instruction stream is fetched from storage for execution. Static scheduling is achieved by moving code and thereby reordering the instruction sequence before execution. This reordering produces an equivalent instruction stream that will more fully utilize the hardware through parallel processing. Such static scheduling is typically done at compile time. However, the reordered instructions remain in their original form and conventional parallel processing still requires some form of dynamic determination just prior to execution of the instructions in order to decide whether to execute the next two instructions serially or in parallel.

20

Such scheduling techniques can improve the overall performance of a pipelined computer, but cannot alone satisfy the ever present demands for increased performance. In that regard, many of the recent proposals for general purpose computing are related to the exploitation of parallelism at the instruction level beyond that attained by pipelining. For example, further instruction level parallelism has been achieved explicitly by issuing multiple instructions per cycle with so-called superscalar machines, rather than implicitly as with dynamic scheduling of single instructions or with vector machines. The name superscalar for machines that issue multiple instructions per cycle is to differentiate them from scalar machines that issue one instruction per cycle.

30

In a typical superscalar machine, the opcodes in a fetched instruction stream are decoded and analyzed dynamically by instruction issue logic in order to determine whether the instructions can be executed in parallel. The

criteria for such last-minute dynamic scheduling are unique to each instruction set architecture, as well for the underlying implementation of that architecture in any given instruction processing unit. Its effectiveness is therefore limited by the complexity of the logic to determine which combinations of instructions can be executed in parallel, and the cycle time of the instruction processing unit is likely to be increased. The increased hardware and cycle time for such superscalar machines become even a bigger problem in architectures which have hundreds of different instructions.

There are other deficiencies with dynamic scheduling, static scheduling, or combinations thereof. For example, it is necessary to review each scalar instruction anew every time it is fetched for execution to determine its capability for parallel execution. There has been no way provided to identify and flag ahead of time those scalar instructions which have parallel execution capabilities.

Another deficiency with dynamic scheduling of the type implemented in super scalar machines is the manner in which scalar instructions are checked for possible parallel processing. Superscalar machines check scalar instructions based on their opcode descriptions, and no way is provided to take into account hardware utilization. Also, instructions are issued in FIFO fashion thereby eliminating the possibility of selective grouping to avoid or minimize the occurrence of interlocks.

There are some existing techniques which do seek to consider the hardware requirements for parallel instruction processing. One such system is a form of static scheduling called the Very Long Instruction Word machine in which a sophisticated compiler rearranges instructions so that hardware instruction scheduling is simplified. In this approach the compiler must be more complex than standard compilers so that a bigger window can be used for purposes of finding more parallelism in an instruction stream. But the resulting instructions may not necessarily be object code compatible with the pre-existing architecture, thereby

solving one problem while creating additional new problems. Also, substantial additional problems arise due to frequent branching which limits its parallelism.

Therefore, none of these prior art approaches to parallel processing have been sufficiently comprehensive to minimize all possible interlocks, while at the same time avoiding major redesign of the architected instruction set and avoiding complex logic circuits for dynamic decoding of fetched instructions.

10 Accordingly, what is needed is an improvement in digital data processing which facilitates the execution of existing machine instructions in parallel in order to increase processor performance. Since the number of instructions executed per second is a product of the basic cycle time of the processor and the average number of cycles required per instruction completion, what is needed is a solution which takes both of these parameters under consideration. More specifically, a mechanism is needed that reduces the number of cycles required for the execution of an instruction for a given architecture. In addition, an improvement is needed which reduces the complexity of the hardware necessary to support parallel instruction execution, thus minimizing any possible increase in cycle time. Additionally, it would be highly desirable for the proposed improvement to provide compatibility of the implementation with an already defined system architecture while introducing parallelism at the instruction level of both new and existing machine code.

20

30 Brief Summary and Objects of the Invention

In view of the foregoing, it is an object of the present invention to provide a method for statically analyzing, at a time prior to instruction decode and execution, a sequence of existing instructions to generate compound instructions formed by adjacent grouping of existing instructions capable of parallel execution. A related object is to add relevant control information to the

instruction stream including grouping information indicating where a compound instruction starts as well as indicating the number of existing instructions which are incorporated into each compound instruction.

Yet another object is to analyze a large window of an instruction byte stream prior to instruction fetch, with the window being adjustable to different positions in the instruction byte stream in order to achieve optimum selective grouping of individual adjacent instructions which form a compound instruction.

10

A further object is to provide an instruction compounding method with the aforementioned characteristics which is applicable to complex instruction architectures having variable length instructions and having data intermixed with instructions, and which is also applicable to RISC architectures wherein instructions are usually a constant length and wherein data is not mixed with instructions.

20

An additional object is to provide a method of pre-processing an instruction stream to create compound instructions, wherein the method can be implemented by software and/or hardware at various points in the computer system prior to instruction decoding and execution. A related object is to provide a method of pre-processing existing instructions which operates on a binary instruction stream as part of a post-compiler, or as part of an in-memory compounder, or as part of cache instruction compounding unit, and which can start compounding instructions at the beginning of a byte stream without knowing the boundaries of the instructions.

30

Thus, the invention contemplates a method of pre-processing an instruction stream to create compound instructions composed of scalar instructions which have still retained their original contents. Compound instructions are created without changing the object code of the scalar instructions which form the compound instruction, thereby allowing existing programs to realize a performance improvement on a compound instruction machine while

maintaining compatibility with previously implemented scalar instruction machines.

10 More specifically, the invention provides a set of compounding rules based on an analysis of existing instructions to separate them into different classes. The analysis determines which instructions qualify, either with instructions in their own class or with instructions in other classes, for parallel execution in a particular hardware configuration. Such compounding rules are used as a standard for pre-processing an instruction stream in order to look for groups of two or more adjacent scalar instructions that can be executed in parallel. In some instances certain types of interlocked instructions can be compounded for parallel execution where the interlocks are collapsible in a particular hardware configuration. In other configurations where the interlocks are non-collapsible, the instructions having data dependent or hardware dependent interlocks are excluded from groups forming compound instructions.

20 Each compound instruction is identified by control information such as tags associated with the compound instruction, and the length of a compound instruction is scalable over a range beginning with a set of two scalar instructions up to whatever maximum number that can be executed in parallel by the specific hardware implementation. Since the compounding rules are based on an identification of classes of instructions rather than on individual instruction, complex matrices showing all possible combinations of specific individual instructions are no longer needed. While keeping their original sequence intact, individual instructions are selectively grouped and combined with one or more other adjacent scalar instructions to form a compound instruction which contains scalar instructions which still have object code compatibility with non-compound scalar instructions. Control information is appended to identify information relevant to the execution of the compound instructions.

30

These and other objects, features and advantages of the

invention will be apparent to those skilled in the art in view of the following detailed description and accompanying drawings.

Brief Description of the Drawings

FIG. 1 is a high level schematic diagram of the invention;

10 FIG. 2 is a timing diagram for a uniprocessor implementation showing the parallel execution of certain non-interlocked instructions which have been selectively grouped in a compound instruction stream;

FIG. 3 is a timing diagram for a multiprocessor implementation showing the parallel execution of scalar and compound instructions which are not interlocked;

FIG. 4 illustrates an example of a possible selective categorization of a portion of the instructions executed by an existing scalar machine;

20 FIG. 5 shows a typical path taken by a program from source code to actual execution;

FIG. 6 is a flow diagram showing generation of a compound instruction set program from an assembly language program;

FIG. 7 is a flow diagram showing execution of a compound instruction set program;

FIGS. 8 is an analytical chart for instruction stream texts with identifiable instruction reference points;

30 FIG. 9 is an analytical chart for an instruction stream text with variable length instructions without a reference point, showing their related sets of possible compound identifier bits;

FIG. 10 illustrates a logical implementation of an instruction compound facility for handling the instruction stream text of FIG. 9;

FIG. 11 is a flow diagram for compounding an instruction stream having reference tags to identify instruction boundary reference points;

FIG. 12 shows an exemplary compound instruction control

field;

FIG. 13 is a flow chart for developing and using compounding rules applicable to a specific computer system hardware configuration and its particular architected instruction set;

FIG. 14 shows how different groupings of valid non-interlocked pairs of instructions form multiple compound instructions for sequential or branch target execution;

10 FIG. 15 shows how different groupings of valid non-interlocked triplets of instructions form multiple compound instructions for sequential or branch target execution;

FIG. 16 is a flow chart for compounding an instruction stream like the one shown in FIG. 9 which includes variable length instructions without boundary reference points; and

FIG. 17 is a chart showing typical compoundable pairs of instruction categories for the portion of the System/370 instruction set shown in FIG. 4.

Detailed Description of Preferred Embodiments

20

The essence of the present invention is the pre-processing of a set of instructions or program to be executed by a computer, to determine statically which non-interlocked instructions may be combined into compound instructions, and the appending of control information to identify such compound instructions. Such determination is based on compounding rules which are developed for an instruction set of a particular architecture. Existing scalar instructions are categorized based on an analysis of their operands, hardware utilization and function, so that

30 grouping of instructions by compounding to avoid non-collapsible interlocks is based on instruction category comparison rather than specific instruction comparison.

As shown in the various drawings and described in more detail hereinafter, this invention called a Scalable Compound Instruction Set Machine (SCISM) provides for a stream of scalar instructions to be compounded or grouped together before instruction decode time so that they are

already flagged and identified for selective simultaneous parallel execution by appropriate instruction execution units. Since such compounding does not change the object code, existing programs can realize a performance improvement while maintaining compatibility with previously implemented systems.

10 As generally shown in Fig. 1, an instruction compounding unit 20 takes a stream of binary scalar instructions 21 (with or without data included therein) and selectively groups some of the adjacent scalar instructions to form encoded compound instructions. A resulting compounded instruction stream 22 therefore combines scalar instructions not capable of parallel execution and compound instructions formed by groups of scalar instructions which are capable of parallel execution. When a scalar instruction is presented to an instruction processing unit 24, it is routed to the appropriate functional unit for serial execution. When a compound instruction is presented to the instruction processing unit 24, its scalar components are each routed to their appropriate functional unit or
20 interlock collapsing unit for simultaneous parallel execution. Typical functional units include but are not limited to an arithmetic and logic unit (ALU) 26, 28, a floating point arithmetic unit (FP) 30, and a store address generation unit (AU) 32.

30 It is to be understood that the technique of the invention is intended to facilitate the parallel issue and execution of instructions in all computer architectures that process multiple instructions per cycle (although certain instructions may require more than one cycle to be executed).

As shown in Fig. 2, the invention can be implemented in a uniprocessor environment where each functional execution unit executes a scalar instruction (S) or alternatively a compounded scalar instruction (CS). As shown in the drawing, an instruction stream 33 containing a sequence of scalar and compounded scalar instructions has control tags (T) associated with each compound instruction. Thus, a

10 first scalar instruction 34 could be executed singly by functional unit A in cycle 1; a triplet compound instruction 36 identified by tag T3 could have its three compounded scalar instructions executed in parallel by functional units A, C and D in cycle 2; another compound instruction 38 identified by tag T2 could have its pair of compounded scalar instructions executed in parallel by functional units A and B in cycle 3; a second scalar instruction 40 could be executed singly by functional unit C in cycle 4; a large group compound instruction 42 could have its four compounded scalar instructions executed in parallel by functional units A-D in cycle 5; and a third scalar instruction 44 could be executed singly by functional unit A in cycle 6.

20 It is important to realize that multiple compound instructions are capable of parallel execution in certain computer system configurations. For example, the invention could be potentially implemented in a multiprocessor environment as shown in Fig. 3 where a compound instruction is treated as a unit for parallel processing by one of the CPUs (central processing units). As shown in the drawing, the same instruction stream 33 could be processed in only two cycles as follows. In a first cycle, a CPU #1 executes the first scalar instruction 34; the functional units of a CPU #2 execute triplet compound instruction 36; and the functional units of a CPU #3 execute the two compounded scalar instructions in compound instruction 38. In a second cycle, the CPU #1 executes the second scalar instruction 40; the functional units of CPU #2 execute the four compounded scalar instructions in compound instruction 42; and a functional unit of CPU #3 executes the third scalar instruction 44.

30 One example of a computer architecture which can be adapted for handling compound instructions is an IBM® System/370¹¹ instruction level architecture in which multiple scalar instructions can be issued for execution in each machine cycle. In this context a machine cycle refers to all the pipeline steps or stages required to execute a

scalar instruction. A scalar instruction operates on operands representing single-valued parameters. When an instruction stream is compounded, adjacent scalar instructions are selectively grouped for the purpose of concurrent or parallel execution.

10 The instruction sets for various IBM System/370 architectures such as System/370, the System/370 extended architecture (370-XA), and the System/370 Enterprise Systems Architecture (370-ESA) are well known. In that regard, reference is given here to the Principles of Operation of the IBM System/370 (publication # GA22-7000-10 1987), and to the Principles of Operation, IBM Enterprise Systems Architecture/370 (publication #SA22-7200-0 1988).

In general, an instruction compounding facility will look for classes of instruction that may be executed in parallel, and ensure that no interlocks between members of a compound instruction exist that cannot be handled by the hardware. When compatible sequences of instructions are found, a compound instruction is created.

20 More specifically, the System/370 instruction set can be broken into categories of instructions that may be executed in parallel in a particular computer system configuration. Instructions within certain of these categories may be combined or compounded with instructions in the same category or with instructions in certain other categories to form a compound instruction. For example, a portion of the System/370 instruction set can be partitioned into the categories illustrated in Fig. 4. The rationale for this categorization is based on the functional requirements of the System/370 instructions and their hardware utilization in a typical computer system configuration. The rest of the System/370 instructions are not considered specifically for compounding in this exemplary embodiment. This does not preclude them from being compounded by the methods and technique of the present invention disclosed herein. It is noted that the hardware structures required for compound instruction execution can be readily controlled by horizontal microcode, allowing for

30

exploitation of parallelism in the remaining instructions not considered for compounding and not included in the categories of Fig. 4, thereby increasing performance.

10 One of the most common sequences in System/370 programs is to execute an instruction of the TM or RX-format COMPARES (C, CH, CL, CLI, CLM), the result of which is used to control the execution of a BRANCH-on-condition type instruction (BC, BCR) which immediately follows. Performance can be improved by executing the COMPARE and the BRANCH instructions in parallel, and this has sometimes been done dynamically in high performance instruction processors. Some difficulty lies in quickly identifying all the various members of the COMPARE class of instructions and all the members of the BRANCH class of instructions in a typical architecture during the instruction decoding process. This is one reason why the superscalar machines usually consider only a small number of specific scalar instructions for possible parallel processing. In contrast, such limited dynamic scheduling based only on a last-minute comparison of
20 two specific instructions is avoided by the invention, because the analysis of all the members of the classes are accomplished ahead of time in order to develop adequate compounding rules for creating a compound instruction which is sure to work.

30 The enormous problem that arises from dynamic scheduling of individual instructions after fetch is shown by realizing that two-way compounding of the fifty seven individual instructions in Fig. 4 produces a 57 X 57 matrix of more than three thousand possible combinations. This is in sharp contrast to the 10 X 10 matrix of Fig. 17 for the same number of instructions considered from the point of view of possible category combinations, as provided by the present invention.

Many classes of instructions may be executed in parallel, depending on how the hardware is designed. In addition to the COMPARE and BRANCH compoundable pairs described above, many other compoundable combinations capable of parallel execution are possible (See Fig. 17),

such as LOADS (category 7) compounded with RR-format instructions (category 1), BRANCHS (categories 3-5) compounded with LOAD ADDRESS (category 8), and the like.

10 In some instances the sequence order will affect the parallel execution capabilities and therefore determine whether two adjacent instructions can be compounded. In that regard, the row headings 45 identify the category of the first instruction in a byte stream and the column headings 47 identify the category of the next instruction which follows the first instruction. For example, BRANCHES (categories 3-5) followed by certain SHIFTS (category 2) are always compoundable 49, while SHIFTS (category 2) followed by BRANCHES (categories 3-5) are only "sometimes" compoundable 51.

20 The "sometimes" status identified as "S" in the chart of Fig. 17 can often be changed to "always" identified as "A" in the chart by adding additional functional hardware units to the computer system configuration. For example, consider a configuration which supports two-way compounding and which has no add-shift collapsing unit, but instead has a conventional ALU and a separate shifter. In other words, there is no interlock collapsing hardware for handling interlocked ADD and SHIFT instructions. Consider the following instruction sequence:

```
AR R1,R2
SRL R3 by D2
```

30 It is clear that this pair of instructions is compoundable for parallel execution. But in some instances it would not be compoundable due to a non-collapsible interlock, as shown in the following instruction sequence:

```
AR R1,R2
SRL R1 by D2
```

So the chart shows that a category 1 instruction (AR) followed by a category 2 instruction (SRL) is sometimes compoundable 53. By including an ALU which collapses certain interlocks such as the add/shift interlock shown above, the S could become an A in the chart of Fig. 17. Accordingly, the compounding rules must be updated to

reflect any changes which are made in the particular computer system configuration.

As an additional example, consider the instructions contained in category 1 compounded with instructions from that same category in the following instruction sequence:

AR R1,R2

SR R3,R4

10 This sequence is free of data hazard interlocks and produces the following results which comprise two independent System/370 instructions:

R1 = R1 + R2

R3 = R3 - R4

20 Executing such a sequence would require two independent and parallel two-to-one ALU's designed to the instruction level architecture. Thus, it will be understood that these two instructions can be grouped to form a compound instruction in a computer system configuration which has two such ALU's. This example of compounding scalar instructions can be generalized to all instruction sequence pairs that are free of data dependent interlocks and also of hardware dependent interlocks.

30 In any actual instruction processor, there will be an upper limit to the number of individual instructions that can comprise a compound instruction. This upper limit must be specifically incorporated into the hardware and/or software unit which is creating the compound instructions, so that compound instructions will not contain more individual instructions (e.g., pair group, triplet group, group of four) than the maximum capability of the underlying execution hardware. This upper limit is strictly a consequence of the hardware implementation in a particular computer system configuration -- it does not restrict either the total number of instructions that may be considered as candidates for compounding or the length of the group window in a given code sequence that may be analyzed for compounding.

In general, the greater the length of a group window being analyzed for compounding, the greater the parallelism

that can be achieved due to more advantageous compounding combinations. In this regard, consider the sequence of instructions in the following Table 1:

Table 1

	X1		;any compoundable instruction
	X2		;any compoundable instruction
10	LOAD	R1,(X)	;load R1 from memory location X
	ADD	R3,R1	;R3 = R3 + R1
	SUB	R1,R2	;R1 = R1 - R2
	COMP	R1,R3	;compare R1 with R3
	X3		;any compoundable instruction
	X4		;any compoundable instruction

If the hardware imposed upper limit on compounding is two (at most, two instructions can be executed in parallel in the same cycle), then there are a number of ways to compound this sequence of instructions depending on the scope of the compounding software.

If the scope of compounding were equal to four, then the compounding software would consider together (X1, X2, LOAD, ADD) and then slide forward one instruction at a time to consider together (X2, LOAD, ADD, SUB) and (LOAD, ADD, SUB, COMP) and (ADD, SUB, COMP, X3) and (SUB, COMP, X3, X4), thereby producing the following optimum pairings as candidates for a compound instruction:

30 [-- X1] [X2 LOAD] [ADD SUB] [COMP X3] [X4 --]

This optimum pairing provided by the invention completely relieves the interlocks between the LOAD and ADD and between the SUB and COMP, and provides the additional possibilities of X1 being compounded with its preceding instruction and of X4 being compounded with its following instruction.

On the other hand, a superscalar machine which pairs instructions dynamically in its instruction issue logic on strictly a FIFO basis, would produce only the following pairings as candidates for parallel execution:

40 [X1 X2] [LOAD ADD] [SUB COMP] [X3 X4]

This inflexible pairing incurs the full penalty of

certain interlocking instructions, and only partial benefits of parallel processing are achieved.

The self explanatory flow chart of Fig. 13 shows the various steps taken to determine which adjacent existing instructions in a byte stream are in categories or classes which qualify them for being grouped together to form a compound instruction for a particular computer system configuration.

10 Referring to Fig. 5, there are many possible locations in a computer system where compounding may occur, both in software and in hardware. Each has unique advantages and disadvantages. As shown in Fig. 5, there are various stages that a program typically takes from source code to actual execution. During the compilation phase, a source program is translated into machine code and stored on a disk 46. During the execution phase the program is read from the disk 46 and loaded into a main memory 48 of a particular computer system configuration 50 where the instructions are executed by appropriate instruction processing units 52, 54, 56. Compounding could take place anywhere along this path. In general as the compounder is located closer to an instruction processing unit or CPUs, the time constraints become more stringent. As the compounder is located further from the CPU, more instructions can be examined in a large sized instruction stream window to determine the best grouping for compounding for increasing execution performance. However such early compounding tends to have more of an impact on the rest of the system design in terms of additional development and cost requirements.

20

30

One of the important objects of the invention is to provide a technique for existing programs written in existing high level languages or existing assembly language programs to be processed by software means which can identify sequences of adjacent instructions capable of parallel execution by individual functional units.

The flow diagram of Fig. 6 shows the generation of a compound instruction set program from an assembly

language program in accordance with a set of customized compounding rules 58 which reflect both the system and hardware architecture. The assembly language program is provided as an input to a software compounding facility 59 that produces the compound instruction program. Successive blocks of instructions having a predetermined length are analyzed by the software compounding facility 59. The length of each block 60, 62, 64 in the byte stream which contains the group of instructions considered together for compounding is dependent on the complexity of the compounding facility.

As shown in Fig. 6, this particular compounding facility is designed to consider two-way compounding for "m" number of fixed length instructions in each block. The primary first step is to consider if the first and second instructions constitute a compoundable pair, and then if the second and third constitute a compoundable pair, and then if the third and fourth constitute a compoundable pair, all the way to the end of the block.

Once the various possible compoundable pairs C1-C5 have been identified, an additional very desirable step is to determine the optimum choice of compound instructions formed by adjacent scalar instructions for parallel execution. In the example shown, the following different sequences of compounded instructions are possible (assuming no branching): I1, C2, I4, I5, C3, C5, I10; I1, C2, I4, I5, I6, C4, I9, I10; C1, I3, I4, I5, C3, C5, I10; C1, I3, I4, I5, I6, C4, I9, I10. Based on the particular hardware configuration, the compounding facility can select the preferred sequence of compounded instructions and use flags or identifier bits to identify the optimum sequence of compound instructions.

If there is no optimum sequence, all of the compoundable adjacent scalar instructions can be identified so that a branch to a target located amongst various compound instructions can exploit any of the compounded pairs which are encountered (See Fig. 14). Where multiple compounding units are available, multiple

successive blocks in the instruction stream could be compounded at the same time.

10 The specific design of a software compounding facility will not be discussed here because the details are unique to a given instruction set architecture and underlying implementation. Although the design of such compounding programs is somewhat similar in concept to modern compilers which perform instruction scheduling and other optimizations based on a specific machine architecture, the criteria used to complete such compounding are unique to this invention, as best shown in the flow chart of Fig. 13. In both instances, given an input program and a description of the instruction set and also of the hardware architectures (i.e., the structural aspects of the implementation), an output program is produced. In the case of the modern compiler, the output is an optimized new sequence of existing instructions. In the case of the invention, the output is a series of compound instructions each formed by a group of adjacent scalar instructions capable of parallel execution, with the compound instructions being intermixed with non-compounded scalar instructions, and with the necessary control bits for execution of the compound instructions included as part of the output.

20

30 Of course, it is easier to pre-process an instruction stream for the purpose of creating compound instructions if known reference points already exist to indicate where instructions begin. As used herein, a reference point means some marking field or other indicator which provides information about the location of instruction boundaries. In many computer systems such a reference point is expressly known only by the compiler at compile time and only by the CPU when instructions are fetched. Such a reference point is unknown between compile time and instruction fetch unless a special reference tagging scheme is adopted.

When compounding is done after compile time, a compiler could indicate with reference tags (see Fig. 11)

which bytes contain the first byte of an instruction and which contain data. This extra information results in a more efficient compounder since exact instruction locations are known. Of course, the compiler could identify instructions and differentiate between instructions and data in other ways in order to provide the compounder with specific information indicating instruction boundaries.

10 When such instruction boundary information is known, the generation of the appropriate compounding identifier bits proceeds in a straightforward manner based on the compounding rules developed for a particular architecture and system hardware configuration (See Fig. 8). When such instruction boundary information is not known, and the instructions are of variable length, a more complex problem is presented (See Figs. 9 and 16). Incidentally, these figures are based on a preferred encoding scheme described in more detail in Table 2A below, wherein two-way compounding provides a tag bit of "1" if an
20 instruction is compounded with the next instruction, and a tag bit of "0" if it is not compounded with the next instruction.

The control bits in a control field added by a compounder contain information relevant to the execution of compound instructions and may contain as little or as much information as is deemed effective for a particular implementation. An exemplary 8-bit control field is shown in Fig. 12. However, only the first control bit is required in the simplest embodiment to indicate the beginning of a compound instruction. The other control bits provide additional optional information relating to the execution of the instructions.
30

In an alternate encoding pattern for compounded instructions applicable to both two-way compounding as well as large group compounding, a first control bit is set to "1" to indicate that the corresponding instruction marks the beginning of a compound instruction. All other members of the compound instruction will have their first

control bit set to "0". On occasion, it will not be possible to combine a given instruction with other instructions, so such a given instruction will appear to be a compound instruction of length one. That is, the first control bit will be set to "1", but the first control bit of the following instruction will also be set to "1". Under this alternate encoding scheme, the decoding hardware will be able to detect how many instructions comprise the compound instruction by monitoring all of the identifier bits for a series of scalar instructions, rather than merely monitoring the identifier bit for the beginning of a compound instruction as in the preferred encoding scheme shown below in Tables 2A-2C.

The flow diagram of Fig. 7 shows a typical implementation for executing a compound instruction set program which has been generated by a hardware preprocessor 66 or a software preprocessor 67. A byte stream having compound instructions flows into a compound instruction (CI) cache 68 that serves as a storage buffer providing fast access to compound instructions. CI issue logic 69 fetches compound instructions from the CI Cache and issues their individual compounded instructions to the appropriate functions units for parallel execution.

It is to be emphasized that compound instruction execution units (CI EU) 71 such as ALU's in a compound instruction computer system are capable of executing either scalar instructions one at a time by themselves or alternatively compounded scalar instructions in parallel with other compounded scalar instructions. Also, such parallel execution can be done in different types of execution units such as ALU's, floating point (FP) units 73, storage address-generation units (AU) 75 or in a plurality of the same type of units (FP1, FP2, etc) in accordance with the computer architecture and the specific computer system configuration. Thus, the hardware configurations which can implement the present invention are scalable up to virtually unlimited numbers

of execution units in order to obtain maximum parallel processing performance. Combining several existing instructions into a single compound instruction allows one or more instruction processing units in a computer system to effectively decode and execute those compounded existing instructions in parallel without the delay that arises in conventional parallel processing computer systems.

10 In the simplest exemplary encoding schemes of this application, minimal compounding information is added to the instruction stream as one bit for every two bytes of text (instructions and data). In general, a tag containing control information can be added to each instruction in the compounded byte stream -- that is, to each non-compounded scalar instruction as well as to each compounded scalar instruction included in a pair, triplet, or larger compounded group. As used herein, identifier bits refers to that part of the tag used specifically to identify and differentiate those
20 compounded scalar instructions forming a compounded group from the remaining non-compounded scalar instructions. Such non-compounded scalar instructions remain in the compound instruction program and when fetched are executed singly.

In a system with all 4-byte instructions aligned on a four byte boundary, one tag is associated with each four bytes of text. Similarly, if instructions can be aligned arbitrarily, a tag is needed for every byte of text.

30 In the illustrated embodiment herein, all System/370 instructions are aligned on a halfword (two-byte) boundary with lengths of either two or four or six bytes, one tag with identifier bits is needed for every halfword. In a small grouping example for compounding pairs of adjacent instructions, an identifier bit "1" indicates that the instruction that begins in the byte under consideration is compounded with the following instruction, while a "0" indicates that the instruction

10 that begins in the byte under consideration is not compounded. The identifier bit associated with halfwords that do not contain the first byte of an instruction is ignored. The identifier bit for the first byte of the second instruction in a compounded pair is also ignored. (However, in some branching situations, these identifier bits are not ignored.) As a result, this encoding procedure for identifier bits means that in the simplest case of two-way compounding, only one bit of information is needed by a CPU during execution to identify a compounded instruction.

20 Where more than two scalar instructions can be grouped together to form a compound instruction, additional identifier bits may be required to provide adequate control information. However, in order to reduce the number of bits required for minimal control information, there is still another alternative format for keeping track of the compounding information. For example, even with large group compounding, it is possible to achieve one bit per instruction with the following encoding: the value "1" means to compound with the next instruction, and the value "0" means to not compound with the next instruction. A compound instruction formed with a group of four individual instructions would have a sequence of compounding identifier bits (1,1,1,0). As with the execution of other compound instructions described herein, compounding identifier bits associated with halfwords which are not instructions and therefore do not have any opcodes are ignored at execution time.

30

Under the preferred encoding scheme described in detail below, the minimum number of identifier bits needed to provide the additional information of indicating the specific number of scalar instructions actually compounded is the logarithm to the base 2 (rounded up to the nearest whole number) of the maximum number of scalar instructions that can be grouped to form a compound instruction. For example, if the maximum is

2039640

EN9-90-020

23

two, then one identifier bit is needed for each compound instruction. If the maximum is three or four, then two identifier bits are needed for each compound instruction. If the maximum is five, six, seven or eight, then three identifier bits are needed for each compound instruction. This encoding scheme is shown below in Tables 2A, 2B and 2C:

EN9-90-020

24

TABLE 2A (maximum of two)

<u>Identifier</u>			
Total #	<u>Bits</u>	<u>Encoded meaning</u>	
<u>Compounded</u>			
	0	This instruction is not compounded with its following instruction	none
10	1	This instruction is compounded with its one following instruction	two

TABLE 2B (maximum of four)

<u>Identifier</u>			
Total #	<u>Bits</u>	<u>Encoded meaning</u>	
<u>Compounded</u>			
	00	This instruction is not compounded with its following instruction	none
20	01	This instruction is compounded with its one following instruction	two
	10	This instruction is compounded with its two following instructions	three
	11	This instruction is compounded with its three following instructions	four
30			

EN9-90-020

25

TABLE 2C (maximum of eight)

<u>Identifier</u>	<u>Total #</u>	<u>Bits</u>	<u>Encoded meaning</u>
		<u>Compounded</u>	
		000	This instruction is not compounded with its following instruction none
10		001	This instruction is compounded with its one following instruction two
		010	This instruction is compounded with its two following instructions three
		011	This instruction is compounded with its three following instructions four
		100	This instruction is compounded with its four following instructions five
20		101	This instruction is compounded with its five following instructions six
		110	This instruction is compounded with its six following instructions seven
		111	This instruction is compounded with its seven following instructions eight

30

It will therefore be understood that each halfword needs a tag, but under this preferred encoding scheme the CPU ignores all but the tag for the first instruction in the

instruction stream being executed. In other words, a byte is examined to determine if it is a compound instruction by checking its identifier bits. If it is not the beginning of a compound instruction, its identifier bits are zero. If the byte is the beginning of a compound instruction containing two scalar instructions, the identifier bits are "1" for the first instruction and "0" for the second instruction. If the byte is the beginning of a compound instruction containing three scalar instructions, the identifier bits are "2" for the first instruction and "1" for the second instruction and "0" for the third instruction. In other words, the identifier bits for each half word identify whether or not this particular byte is the beginning of a compound instruction while at the same time indicating the number of instructions which make up the compounded group.

These exemplary methods of encoding compound instructions assume that if three instructions are compounded to form a triplet group, the second and third instructions are also compounded to form a pair group. In other words, if a branch to the second instruction in a triplet group occurs, the identifier bit "1" for the second instruction indicates that the second and third instruction will execute as a compounded pair in parallel, even though the first instruction in the triplet group was not executed.

Of course, the invention is not limited to this particular preferred encoding scheme. Various other encoding rules, such as the alternate encoding scheme previously described, are possible within the scope and teachings of the invention.

It will be apparent to those skilled in the art that the present invention requires an instruction stream to be compounded only once for a particular computer system configuration, and thereafter any fetch of compounded instructions will also cause a fetch of the identifier bits associated therewith. This avoids the need for the

inefficient last-minute determination and selection of certain scalar instructions for parallel execution that repeatedly occurs every time the same or different instructions are fetched for execution in the so-called super scalar machine.

10 Despite all of the advantages of compounding a binary instruction stream, it becomes difficult to do so under certain computer architectures unless a technique is developed for determining instruction boundaries in a byte string. Such a determination is complicated when variable length instructions are allowed, and is further complicated when data and instructions can be intermixed, and when modifications are allowed to be made directly to the instruction stream. Of course, at execution time instruction boundaries must be known to allow proper execution. But since compounding is preferably done a sufficient time prior to instruction execution, a unique technique has been developed to compound instructions without knowledge of where instructions start and without
20 knowledge of which bytes are data. This technique is described generally below and can be used for creating compound instructions formed from adjacent pairs of scalar instructions as well as for creating compound instructions formed from larger groups of scalar instructions. This technique is applicable to all instruction sets of the various conventional types of architectures, including the RISC (Reduced Instruction Set Computers) architectures in which instructions are usually a constant length and are not intermixed with
30 data.

 Generally speaking, the compounding technique provides for the compounding two or more scalar instructions from an instruction stream without knowing the starting point or length of each individual instruction. Typical instructions already include an opcode at a predetermined field location which identifies the instruction and its length. Those adjacent instructions which qualify for parallel execution in a

particular computer system configuration are provided with appropriate tags to indicate they are candidates for compounding. In IBM System/370 architecture where instructions are either two, four or six bytes in length, the field positions for the opcode are presumed based on an estimated instruction length code. The value of each tag based on a presumed opcode is recorded, and the instruction length code in the presumed opcode is used to locate a complete sequence of possible instructions. Once an actual instruction boundary is found, the corresponding correct tag values are used to identify the commencement of a compound instruction, and other incorrectly generated tags are ignored.

This unique compounding technique is exemplified in the drawings of Figs. 8-9 and 14-15 wherein the compounding rules are defined to provide that all instructions which are 2 bytes or 4 bytes long are compoundable with each other (i.e., a 2 byte instruction is capable of parallel execution in this particular computer configuration with another 2 byte or another 4 byte instruction). The exemplary compounding rules further provide that all instructions which are 6 bytes long are not compoundable at all (i.e., a 6 byte instruction is only capable of execution singly by itself in this particular computer configuration). Of course, the invention is not limited to these exemplary compounding rules, but is applicable to any set of compounding rules which define the criteria for parallel execution of existing instructions in a specific configuration for a given computer architecture.

The instruction set used in these exemplary compounding techniques of the invention is taken from the System/370 architecture. By examining the opcode for each instruction, the type and length of each instruction can be determined and the control tag containing identifier bits is then generated for that specific instruction, as described in more detail hereinafter. Of course, the present invention is not limited to any specific

architecture or instruction set, and the aforementioned compounding rules are by way of example only.

The preferred encoding scheme for compound instructions in these illustrated embodiments has already been shown above in Table 2A-2C.

10 In a First case with fixed length instructions having no data intermixed and with a known reference point location for the opcode, the compounding can proceed in accordance with the applicable rules for that particular computer configuration. Since the field reserved for the opcode also contains the instruction length, a sequence of scalar instructions is readily determined, and each instruction in the sequence can be considered as possible candidates for parallel execution with a following instruction. A first encoded value in the control tag indicates the instruction is not compoundable with the next instruction, while a second encoded value in the control tag indicates the instruction is compoundable for parallel execution with the next instruction.

20

In a Second case with variable length instructions having no data intermixed, and with a known reference point location for the opcode and also for the instruction length code (which in System/370 is included as part of the opcode), the compounding can proceed in a routine manner. As shown in Fig. 8, the opcodes indicate an instruction sequence 70 as follows: the first instruction is 6 bytes long, the second and third are each 2 bytes long, the fourth is 4 bytes long, the fifth is 2 bytes long, the sixth is 6 bytes long, and the seventh and eighth are each 2 bytes long.

30

A C-vector 72 in Fig. 8 shows the values for the identifier bits (called compounding bits in the drawings) for this particular sequence 70 of instructions where a reference point indicating the beginning of the first instruction is known. Based on the values of such identifier bits, the second and third instructions form a compounded pair as indicated by the "1" in the identifier

EN9-90-020

30

bit for the second instruction. The fourth and fifth instructions form another compounded pair as indicated by the "1" in the identifier bit for the fourth instruction. The seventh and eighth instructions also form a compounded pair as indicated by the "1" in the identifier bit for the seventh instruction.

10 The C-vector 72 of Fig. 8 is relatively easy to generate when there are no data bytes intermixed with the instruction bytes, and where the instructions are all of the same length with known boundaries.

Another situation is presented in a Third case where instructions are mixed with non-instructions, with a reference point still being provided to indicate the beginning of an instruction. The schematic diagram of Fig. 11 shows one way of indicating an instruction reference point, where every halfword has been flagged with a reference tag to indicate whether or not it contains the first byte of an instruction. This could occur with both fixed length and variable length instructions. By providing the reference point, it is unnecessary to evaluate the data portion of the byte stream for possible compounding. Accordingly, the compounding unit can skip over and ignore all of the non-instruction bytes.

20

A more complicated situation arises where a byte stream includes variable length instructions (without data), but it is not known where a first instruction begins. Since the maximum length instruction is six bytes, and since instructions are aligned on two byte boundaries, there are three possible starting points for the first instruction the the stream. Accordingly, the technique provides for considering all possible starting points for the first instruction in the text of a byte stream 79, as shown in Fig. 9.

30

Sequence 1 assumes that the first instruction starts with the first byte, and proceeds with compounding on that premise. In this exemplary embodiment, the length field is also determinative of the C-vector value for

each possible instruction. Therefore a C-vector 74 for Sequence 1 only has a "1" value for the first instruction of a possible compounded pair formed by combinations of 2 byte and 4 byte instructions.

10 Sequence 2 assumes that the first instruction starts at the third byte (the beginning of the second halfword), and proceeds on that premise. The value in the length field for the third byte is 2 indicating the next instruction begins with the fifth byte. By proceeding through each possible instruction based on the length field value in the preceding instruction, the entire potential instructions of Sequence 2 are generated along with the possible identifier bits as shown in a C-vector 76.

20 Sequence 3 assumes that the first instruction starts at the fifth byte (the beginning of the third halfword), and proceeds on that premise. The value in the length field for the fifth byte is 4 indicating the next instruction begins with the ninth byte. By proceeding through each possible instruction based on the length field value in the preceding instruction, the entire potential instructions of Sequence 3 are generated along with the possible identifier bits as shown in a C-vector 78.

30 In some instances the three different Sequences of potential instructions will converge into one unique sequence. In Fig. 9 it is noted that the three Sequences converge on instruction boundaries at the end 80 of the eighth byte. Sequences 2 and 3, while converging on instruction boundaries at the end 82 of the fourth byte, are out-of-phase in compounding until the end of the sixteenth byte. In other words, the two sequences consider different pairs of instructions based on the same sequence of instructions. Since the seventeenth byte begins a non-compoundable instruction at 84, the out-of-phase convergence is ended.

When no valid convergence occurs, it is necessary to continue all three possible instruction sequences to the

EN9-90-020

32

end of the window. However, where valid convergence occurs and is detected, the number of sequences collapses from three to two (one of the identical sequences becomes inoperative), and in some instances from two to one.

10 Thus, prior to convergence, tentative instruction boundaries are determined for each possible instruction sequence and identifier bits assigned for each such instruction indicating the location of the potential compound instructions. It is apparent from Fig. 9 that this technique generates three separate identifier bits for every two text bytes. In order to provide consistency with the pre-processing done in the aforementioned first, second and third cases, it is desirable to reduce the three possible sequences to a single sequence of identifier bits where only one bit is associated with each halfword. Since the only information needed is whether the current instruction is compounded with the following instruction, the three bits can be logically ORed to produce a single sequence in a
20 CC-vector 86.

For purposes of parallel execution, the composite identifier bits of a composite CC-vector are equivalent to the separate C-vectors of the individual three Sequences 1-3. In other words, the composite identifier bits in the CC-vector allow any of the three possible sequences to execute properly in parallel for compound instructions or singly for non-compounded instructions. The composite identifier bits also work properly for branching. For example, if a branch to the beginning 88
30 of the ninth byte occurs, then the ninth byte must begin an instruction. Otherwise there is an error in the program. The identifier bit "1" associated with the ninth byte is used and correct parallel execution of such instruction with its next instruction proceeds.

The various steps in the compounding method shown in Fig. 9 as described above are illustrated in the self-explanatory flow chart of Fig. 16.

The best time for providing reference point

information for instruction boundaries is at the time of compiling. Reference tags 101 could be added at compile time to identify the beginning of each instruction, as shown in Fig. 11. This enables the compounder to proceed with the simplified technique for the aforementioned First, Second and Third cases. Of course, the compiler could identify instruction boundaries and differentiate between instructions and data in other ways, in order to simplify the work of the compounding unit and avoid the complications of a technique like the one shown in Fig. 9.

Fig. 10 shows a flow diagram of a possible implementation of a compounder for handling instruction streams like the one in Fig. 9. A multiple number of compounder units 104, 106, 108 are shown, and for efficiency purposes this number could be as large as the number of halfwords that could be held in a text buffer. In this version, the three compounder units would begin their processing sequences at the first, third, and fifth bytes, respectively. Upon finishing with a possible instruction sequence, each compounder starts examining the next possible sequence offset by six bytes from its previous sequence. Each compounder produces compound identifier bits (C-vector values) for each halfword in the text. The three sequences from the three compounders are ORed 110 and the resulting composite identifier bits (CC-vector values) are stored in association with their corresponding textual bytes.

One beneficial advantage provided by the composite identifier bits in the CC-vector is the creation of multiple valid compounding bit sequences based on which instruction is addressed by a branch target. As best shown in Figs. 14-15, differently formed compounded instructions are possible from the same byte stream.

Fig. 14 shows the possible combinations of compounded instructions when the computer configuration provides for parallel issuance and execution of no more than two instructions. Where an instruction stream 90

containing compounded instructions is processed in normal sequence, the Compound Instruction I will be issued for parallel execution based on decoding of the identifier bit for the first byte in a CC-vector 92. However, if a branch to the fifth byte occurs, the Compound Instruction II will be issued for parallel execution based on decoding of the identifier bit for the fifth byte.

10 Similarly, a normal sequential processing of another compounded byte stream 94 will result in Compound Instructions IV, VI and VIII being sequentially executed (the component instructions in each compound instruction being executed in parallel). In contrast, branching to the third byte in the compounded byte stream will result in Compound Instructions V and VII being sequentially executed, and the instruction beginning at the fifteenth byte (it forms the second part of Compound Instruction VIII) will be issued and executed singly, all based in the identifier bits in the CC-vector 96 .

20 Branching to the seventh byte will result in Compound Instructions VI and VIII being sequentially executed, and branching to the eleventh byte will result in Compound Instruction VIII being executed. In contrast, branching to the ninth byte in the compounded byte stream will result in Compound Instruction VII being executed (it is formed by the second part of Compound Instruction VI and the first part of Compound Instruction VIII).

30 Thus, the identifier bits "1" in the CC-vector 96 for Compound Instructions IV, VI and VIII are ignored when either of the Compound Instructions V or VII is being executed. Alternatively the identifier bits "1" in the CC-vector 96 for Compound Instructions V and VII are ignored when any of Compound Instructions IV, VI or VIII are executed.

Fig. 15 shows the possible combinations of compounded instructions when the computer configuration provides for parallel issuance and execution of up to three instructions. Where an instruction stream 98

containing compounded instructions is processed in normal sequence, the Compound Instructions X (a triplet group) and XIII (a pair group) will be executed. In contrast, branching to the eleventh byte will result in Compound Instruction XI (a triplet group) being executed, and branching to the thirteenth byte will result in Compound Instruction XII (a different triplet group) being executed.

10 Thus, the identifier bits "2" in a CC-vector 99 for Compound Instructions XI and XII are ignored when Compound Instructions X and XIII are executed. On the other hand when Compound Instruction XI is executed, the identifier bits for the other three Compound Instructions X, XII, XIII are ignored. Similarly when Compound Instruction XII is executed, the identifier bits for the other three Compound Instructions X, XI, XIII are ignored.

20 There are many possible designs for an instruction compounding unit depending on its location and the knowledge of the text contents. In the simplest situation, it would be desirable for a compiler to indicate with reference tags which bytes contain the first byte of an instruction and which contain data. This extra information results in a more efficient compounder since exact instruction locations are known. This means that compounding could always be handled as in the First, Second and Third case situations in order to generate the C-vector identifier bits for each compound instruction. A compiler could also add other information
30 such as static branch prediction or even insert directives to the compounder.

Other ways could be used to differentiate data from instructions where the instruction stream to be compounded is stored in memory. For example, if the data portions are infrequent, a simple list of addresses containing data would require less space than reference tags. Such combinations of a compounder in hardware and software provide many options for efficiently producing

2039640

EN9-90-020

36

compound instructions.

While exemplary preferred embodiments of the invention have been disclosed, it will be appreciated by those skilled in the art that various modifications and changes can be made without departing from the spirit and scope of the invention as defined by the following claims.

The embodiments of the invention in which an exclusive property or privilege is claimed are defined as follows:

1. A method of obtaining parallel execution of existing instructions in a particular configuration of a data processing system, including:

allocating certain of the existing instructions into multiple categories;

comparing the categories of adjacent existing instructions in an instruction stream to determine whether the adjacent existing instructions are capable of parallel execution in the particular configuration of a data processing system;

identifying with a compounding indicator those adjacent existing instructions which are determined by said comparing step to be capable of parallel execution; and

said allocating and comparing steps being performed prior to the time the existing instructions are fetched for execution.

2. The method of claim 1 wherein said allocating step includes allocating certain instructions into at least one category such that said comparing step causes said identifying step to identify with a compounding indicator at least two adjacent existing instructions from said at least one category for parallel execution.

3. The method of claim 1 wherein said allocating step includes allocating certain instructions into at least two different non-overlapping categories such that said comparing step causes said identifying step to identify with a compounding indicator at least two adjacent existing instructions from said at least two categories, respectively, for parallel execution.

4. The method of claim 1, 2 or 3 wherein said identifying step includes leaving the object code of the existing instructions in its original form for execution singly or for execution in parallel with another instruction.

5. The method of any one of claims 1 to 4 wherein said allocating step includes taking into account data dependent interlocks between instructions as well as the existence of related interlock collapsing functional units in the particular configuration of the data processing system.

6. The method of any one of claims 1 to 5 wherein said allocating step includes taking into account hardware dependent interlocks between instructions as well as the existence of related interlock collapsing functional units in the particular configuration of the data processing system.

7. The method of any one of claims 1 to 6 wherein said comparing step includes comparing a first existing instruction with a second adjacent following instruction for possible compounding with each other, and then comparing the second instruction with a third adjacent following instruction for possible compounding with each other, to identify multiple compound instructions identified by multiple compounding indicators, respectively.

8. The method of claim 7 further including the additional step of determining an optimum sequence of multiple compound instructions for a given portion of an instruction stream.

9. A system for processing a sequence of instructions taken from an existing set of machine instructions in order to prepare the sequence for parallel execution by a particular computer configuration, including:

a set of rules based on the particular computer configuration indicating certain machine instructions which are capable of parallel execution with other machine instructions;

compounding means for pre-processing prior to instruction fetch a binary instruction stream containing a sequence of existing instructions, said compounding means having access to said set of rules for identifying adjacent instructions capable of parallel execution in the particular computer configuration and for producing a compounded sequence of instructions having a control field identifying individual adjacent instructions capable of parallel execution and which together form a compound instruction; and

said compounding means comprising a software pre-processing program used as part of a compiler or a post-compiler.

10. The system of claim 9, wherein said compounding means is a hardware pre-processing unit which processes instructions prior to their being fetched from a cache buffer.

11. The system of claim 9, wherein said compounding means is capable of producing a compounded sequence of instructions having a control field identifying pairs of adjacent instructions capable of parallel execution.
12. The system of claim 9, 10 or 11, wherein said compounding means is capable of producing a compounded sequence of instructions having a control field identifying groups of three or more adjacent instructions capable of parallel execution.
13. The system of any one of claims 9 to 12, wherein said control field comprises at least one identifier bit associated with each compound instruction.
14. The system of any one of claims 9 to 13, wherein said control field includes an identifier bit portion and a control bit portion different from said identifier bit portion.
15. The system of any one of claims 9 to 14, wherein said control field comprises at least one identifier bit associated with each individual instruction forming a compound instruction.
16. The system of any one of claims 9 to 15, wherein said control field comprises at least one identifier bit associated with each individual instruction not forming a compound instruction.
17. The system of any one of claims 9 to 16 wherein said set of rules takes into account data dependent interlocks between instructions as well as the existence of related interlock collapsing functional units in the particular configuration of the data processing system.
18. The system of any one of claims 9 to 17 wherein said set of rules takes into account hardware dependent interlocks between instructions as well as the existence of related interlock collapsing functional units in the particular configuration of the data processing system.

19. A computer implemented method for processing a sequence of binary encoded scalar instructions, each instruction including an operation code, prior to fetching for execution in a data processing system, which places constraints on parallel instruction execution, including the steps of:

transferring said sequence of binary encoded scalar instructions from a data storage unit to an instruction compounding unit;

assigning each operation code in said sequence of instructions to one of a plurality of categories based on a function performed by said data processing system in response to said operation code, the number of said categories being less than the number of operation codes in said sequence of instructions;

storing in said data processing system matrix data that encodes which instruction pairs in said sequence of instructions are compoundable based on the category assigned to the operation code of each adjacent instruction of said instruction pairs; and

processing groups of instructions in said compounding unit to generate a compounding instruction for said instruction pairs by comparing to said matrix data the category assigned in said assigning step respectively to each instruction of said instruction pairs in said groups of instructions.

20. The method of claim 19 wherein said processing step is performed prior to the time the existing instructions are fetched for execution.

21. The method of claim 19 or 20 wherein said processing step retains the object code of the existing instructions in its original form for execution singly or for execution in parallel with another instruction when a fetched instruction has a compounding indicator that its associated instruction can be executed in parallel with an adjacent instruction of a fetched instruction stream.

22. The method of claim 19, 20 or 21 wherein said processing step takes into account data dependent interlocks between instructions as well as the existence of related interlock collapsing functional units in the particular configuration of a data processing system.

23. The method of any one of claims 19 to 22 wherein said processing step takes into account hardware dependent interlocks between instructions as well as the existence of related interlock collapsing functional units in the particular configuration of a data processing system.

24. The method of any one of claims 19 to 23 wherein said processing step compares a first existing instruction with a second adjacent instruction for possible compounding with each other, and then compares the second instruction with a third adjacent following instruction for possible compounding with each other, to identify multiple compound instructions identified by multiple compounding indicators, respectively.

25. The method of claim 24 further including the additional step of determining an optimum sequence of multiple compound instructions for a given portion of an instruction stream.

26. The method of any one of claims 19 to 25 wherein said compounding instruction is associated with each instruction of the instruction stream which can be grouped as adjacent instructions which can be executed in parallel in a particular configuration of a data processing systems has a value which indicates the number of instructions following the instruction in the instruction stream which can be compounded with the instruction upon execution of the instruction when the original instruction is fetched with the associated compounding indicator for execution in a particular computer configuration.

27. A computer system for processing a sequence of binary encoded scalar instructions, each instruction including an operation code, prior to fetching for execution in a data processing system, which places constraints on parallel instruction execution, including the steps of:

means for transferring said sequence of binary encoded scalar instructions from a data storage unit to an instruction compounding unit;

means for assigning each operation code in said sequence of instructions to one of a plurality of categories based on a function performed by said data processing system in response to said operation code, the number of said categories being less than the number of operation codes in said sequence of instructions;

means for storing in said data processing system matrix data that encodes which instruction pairs in said sequence of instructions are compoundable based on the category assigned to the operation code of each adjacent instruction of said instruction pairs; and

means for processing groups of instructions in said compounding unit to generate a compounding instruction for said instruction pairs by comparing to said matrix data the category assigned in said means for assigning respectively to each instruction of said instruction pairs in said groups of instructions.

28. The system of claim 27 wherein said means for processing performs its processing prior to the time the existing instructions are fetched for execution.

29. The system of claim 27 or 28 wherein said means for processing retains the object code of the existing instructions in its original form for execution singly or for execution in parallel with another instruction when a fetched instruction has a compounding indicator that its associated instruction can be executed in parallel with an adjacent instruction of a fetched instruction stream.

30. The system of claim 27, 28 or 29 wherein said means for processing takes into account data dependent interlocks between instructions as well as the existence of related interlock collapsing functional units in the particular configuration of a data processing system.

31. The system of any one of claims 27 to 30 wherein said means for processing takes into account hardware dependent interlocks between instructions as well as the existence of related interlock collapsing functional units in the particular configuration of a data processing system.

32. The system of any one of claims 27 to 31 wherein said means for processing compares a first existing instruction with a second adjacent instruction for possible compounding with each other, and then compares the second instruction with a third adjacent following instruction for possible compounding with each other, to identify multiple compound instructions identified by multiple compounding indicators, respectively.

33. The system of claim 32 further including means for determining an optimum sequence of multiple compound instructions for a given portion of an instruction stream.

34. The system of any one of claims 27 to 33 wherein said compounding instruction is associated with each instruction of the instruction stream which can be grouped as adjacent instructions which can be executed in parallel in a particular configuration of a data processing systems has a value which indicates the number of instructions following the instruction in the instruction stream which can be compounded with the instruction upon execution of the instruction when the original instruction is fetched with the associated compounding indicator for execution in a particular computer configuration.

35. A computer implemented system for processing a sequence of existing binary encoded scalar instructions stored in a data storage unit, each instruction including an operation code, prior to fetching said sequence for execution in a data processing system that is capable of parallel instruction execution but places constraints on said parallel instruction execution, comprising in combination:

an instruction compounding unit;

means for transferring said sequence of binary encoded scalar instructions from said data storage unit to said instruction compounding unit;

said instruction compounding unit assigning each operation code in said sequence of instructions to one of a plurality of categories based on a function performed by said data processing system in response to said operation code, the number of said categories being less than the number of operation codes in said sequence of instructions;

said instruction compounding unit storing data that encodes those instruction pairs in said sequence of instructions that are compoundable based on the category assigned to the operation code of each adjacent instruction of said instruction pairs;

said instruction compounding unit processing groups of instructions to generate a compounding instruction for instruction pairs by comparing in a relational matrix a category assigned respectively to each instruction of the instruction pairs in said groups of instructions.

36. The system of claim 35 wherein said compounding unit is a software pre-processing program used as part of a compiler or post-compiler.

37. The system of claim 35 wherein said compounding unit is a hardware pre-processing unit which processes instructions prior to their being fetched from a cache buffer.

38. The system of claim 35, 36, or 37 wherein said compounding unit is capable of producing a compounded sequence of instructions having a control field identifying pairs of adjacent instructions capable of parallel execution.


39. The system of any one of claims 35 to 38 wherein said compounding unit is capable of producing a compounded sequence of instructions having a control field identifying groups of three or more adjacent instructions capable of parallel execution.

40. The system of claim 38 or 39 wherein said control field comprises at least one identifier bit associated with each compound instruction.

41. The system of claim 38, 39 or 40 wherein said control field includes an identifier bit portion and a control bit portion different from said identifier bit portion.

42. The system of any one of claims 38 to 41 wherein said control field comprises at least one identifier bit associated with each individual instruction forming a compound instruction.

43. The system of any one of claim 38 to 42 wherein said control field comprises at least one identifier bit associated with each individual instruction not forming a compound instruction.



[Signature]
PATENT AGENT





FIG. 2

Richard J. ...
 PATENT AGENT

CYCLE 1 CYCLE 2 INSTRUCTION CYCLES

	FU-A		
	FU-B		
CPU #1	FU-C		
	FU-D	S	S
CPU #2	FU-A	CS	CS
	FU-B	CS	CS
	FU-C	CS	CS
	FU-D	—	CS
CPU #3	FU-A	CS	—
	FU-B	CS	S
	FU-C	—	—
	FU-D	—	—

CENTRAL PROCESSING UNIT = CPU
 SCALAR INSTRUCTION = S
 COMPOUNDED SCALAR INSTRUCTION = CS
 COMPOUND INSTRUCTION = C
 COMPOUND INSTRUCTION TAG = T
 FUNCTIONAL UNIT = FU




MULTI-PROCESSOR IMPLEMENTATION

FIG. 3


Alonzo J. ...
 PATENT AGENT

CATEGORY NUMBER	CATEGORY NAME
1.	RR-FORMAT LOADS, LOGICALS, ARITHMETICS, COMPARES <ul style="list-style-type: none"> • LCR - LOAD COMPLEMENT • LPR - LOAD POSITIVE • LNR - LOAD NEGATIVE • LR - LOAD REGISTER • LTR - LOAD AND TEST • NR - AND • OR - OR • XR - EXCLUSIVE OR • AR - ADD • SR - SUBTRACT • ALR - ADD LOGICAL • SLR - SUBTRACT LOGICAL • CLR - COMPARE LOGICAL • CR - COMPARE
2.	RS-FORMAT SHIFTS (NO STORAGE ACCESS) <ul style="list-style-type: none"> • SRL - SHIFT RIGHT LOGICAL • SLL - SHIFT LEFT LOGICAL • SRA - SHIFT RIGHT ARITHMETIC • SLA - SHIFT LEFT ARITHMETIC • SRDL - SHIFT RIGHT LOGICAL • SLDL - SHIFT LEFT LOGICAL • SRDA - SHIFT RIGHT ARITHMETIC • SLDA - SHIFT LEFT ARITHMETIC
3.	BRANCHES-ON COUNT AND INDEX <ul style="list-style-type: none"> • BCT - BRANCH ON COUNT (RX-FORMAT) • BCTR - BRANCH ON COUNT (RR-FORMAT) • BXH - BRANCH ON INDEX HIGH (RS-FORMAT) • BXLE - BRANCH ON INDEX LOW (RS-FORMAT)
4.	BRANCHES-ON CONDITION <ul style="list-style-type: none"> • BC - BRANCH ON CONDITION (RX-FORMAT) • BCR - BRANCH ON CONDITION (RR-FORMAT)

FIG.
4AFIG.
4BFIG.4FIG.4A


PATENT AGENT

CATEGORY NUMBER	CATEGORY NAME
5.	BRANCHES--AND LINK <ul style="list-style-type: none"> • BAL - BRANCH AND LINK (RX-FORMAT) • BALR - BRANCH AND LINK (RR-FORMAT) • BAS - BRANCH AND SAVE (RX-FORMAT) • BASR - BRANCH AND SAVE (RR-FORMAT)
6.	STORES <ul style="list-style-type: none"> • STCM - STORE CHARACTERS UNDER MASK (0-4-BYTE STORE, RS-FORMAT) • MVI - MOVE IMMEDIATE (ONE BYTE, SI-FORMAT) • ST - STORE (4 BYTES) • STC - STORE CHARACTER (ONE BYTE) • STH - STORE HALF (2 BYTES)
7.	LOADS <ul style="list-style-type: none"> • LH - LOAD HALF (2 BYTES) • L - LOAD (4 BYTES)
8.	LA-LOAD ADDRESS
9.	RX/SI/RS-FORMAT ARITHMETICS, LOGICALS, INSERTS, COMPARES <ul style="list-style-type: none"> • A - ADD • AH - ADD HALF • AL - ADD LOGICAL • N - AND • O - OR • S - SUBTRACT • SH - SUBTRACT HALF • SL - SUBTRACT LOGICAL • X - EXCLUSIVE OR • IC - INSERT CHARACTER • ICM - INSERT CHARACTERS UNDER MASK (0-TO 4-BYTE FETCH) • C - COMPARE • CH - COMPARE HALF • CL - COMPARE LOGICAL • CLI - COMPARE LOGICAL IMMEDIATE • CLM - COMPARE LOGICAL CHARACTER UNDER MASK
10.	TM-TEST UNDER MASK

FIG.4B

PATENT AGENT





FIG. 5

PATENT AGENT



Alfred J. ...
PATENT AGENT




FIG.7

PATENT AGENT

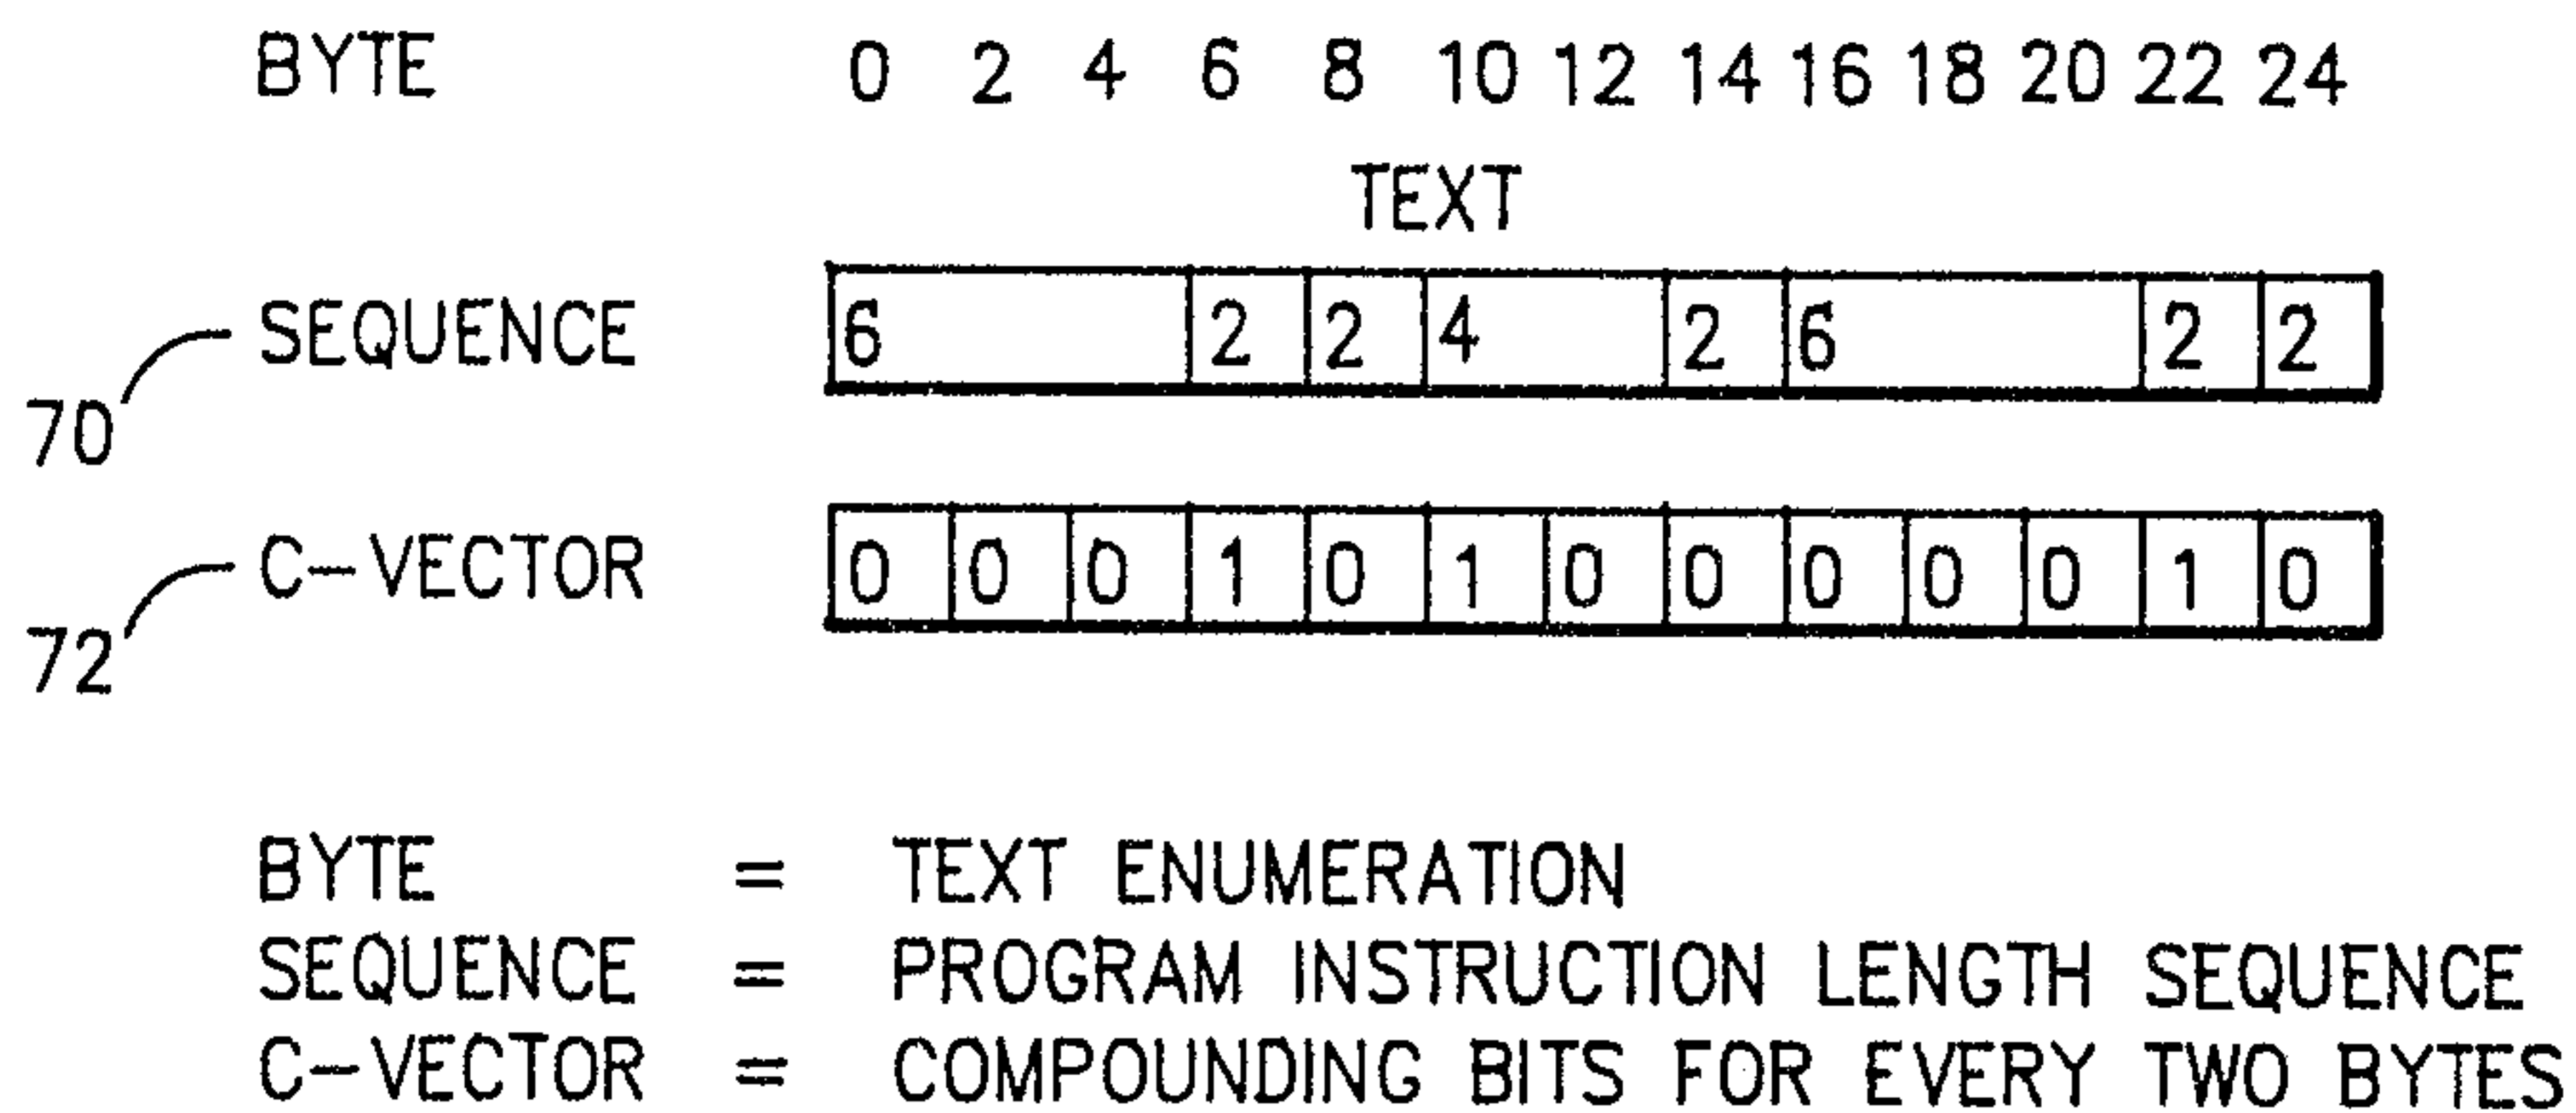


FIG.8





FIG.11

[Handwritten Signature]

PATENT AGENT



LENGTH FIELD = POSSIBLE INSTRUCTION LENGTH CODE FOR EACH TWO BYTES
 BYTE = TEXT ENUMERATION
 SEQUENCE = POTENTIAL PROGRAM INSTRUCTION LENGTH SEQUENCE
 C-VECTOR = POSSIBLE COMPOUNDING BITS FOR EVERY TWO BYTES
 CC-VECTOR = COMPOSITE COMPOUNDING BITS FOR EVERY TWO BYTES

FIG.9

PATENT AGENT




FIG.10

Alvin J. [Signature]
 PATENT AGENT

t ₀	t ₁	t ₂	t ₃	t ₄	t ₅	t ₆	t ₇
----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------

BIT	FUNCTION
t ₀	IF 1, THIS INSTRUCTION MARKS THE BEGINNING OF A COMPOUND INSTRUCTION.
t ₁	IF 1, THEN EXECUTE TWO COMPOUND INSTRUCTIONS IN PARALLEL.
t ₂	IF 1, THEN THIS COMPOUND INSTRUCTION HAS MORE THAN ONE EXECUTION CYCLE.
t ₃	IF 1, THEN SUSPEND PIPELINING.
t ₄	IF INSTRUCTION IS A BRANCH, THEN IF THIS BIT IS 1, THE BRANCH IS PREDICTED TO BE TAKEN.
t ₅	IF 1, THEN THIS INSTRUCTION HAS A STORAGE INTERLOCK FROM A PREVIOUS COMPOUND INSTRUCTION.
t ₆	IF 1, THEN ENABLE DYNAMIC INSTRUCTION ISSUE.
t ₇	IF 1, THEN THIS INSTRUCTION USES AN ALU.

FIG.12

Richard J. ...
 PATENT AGENT

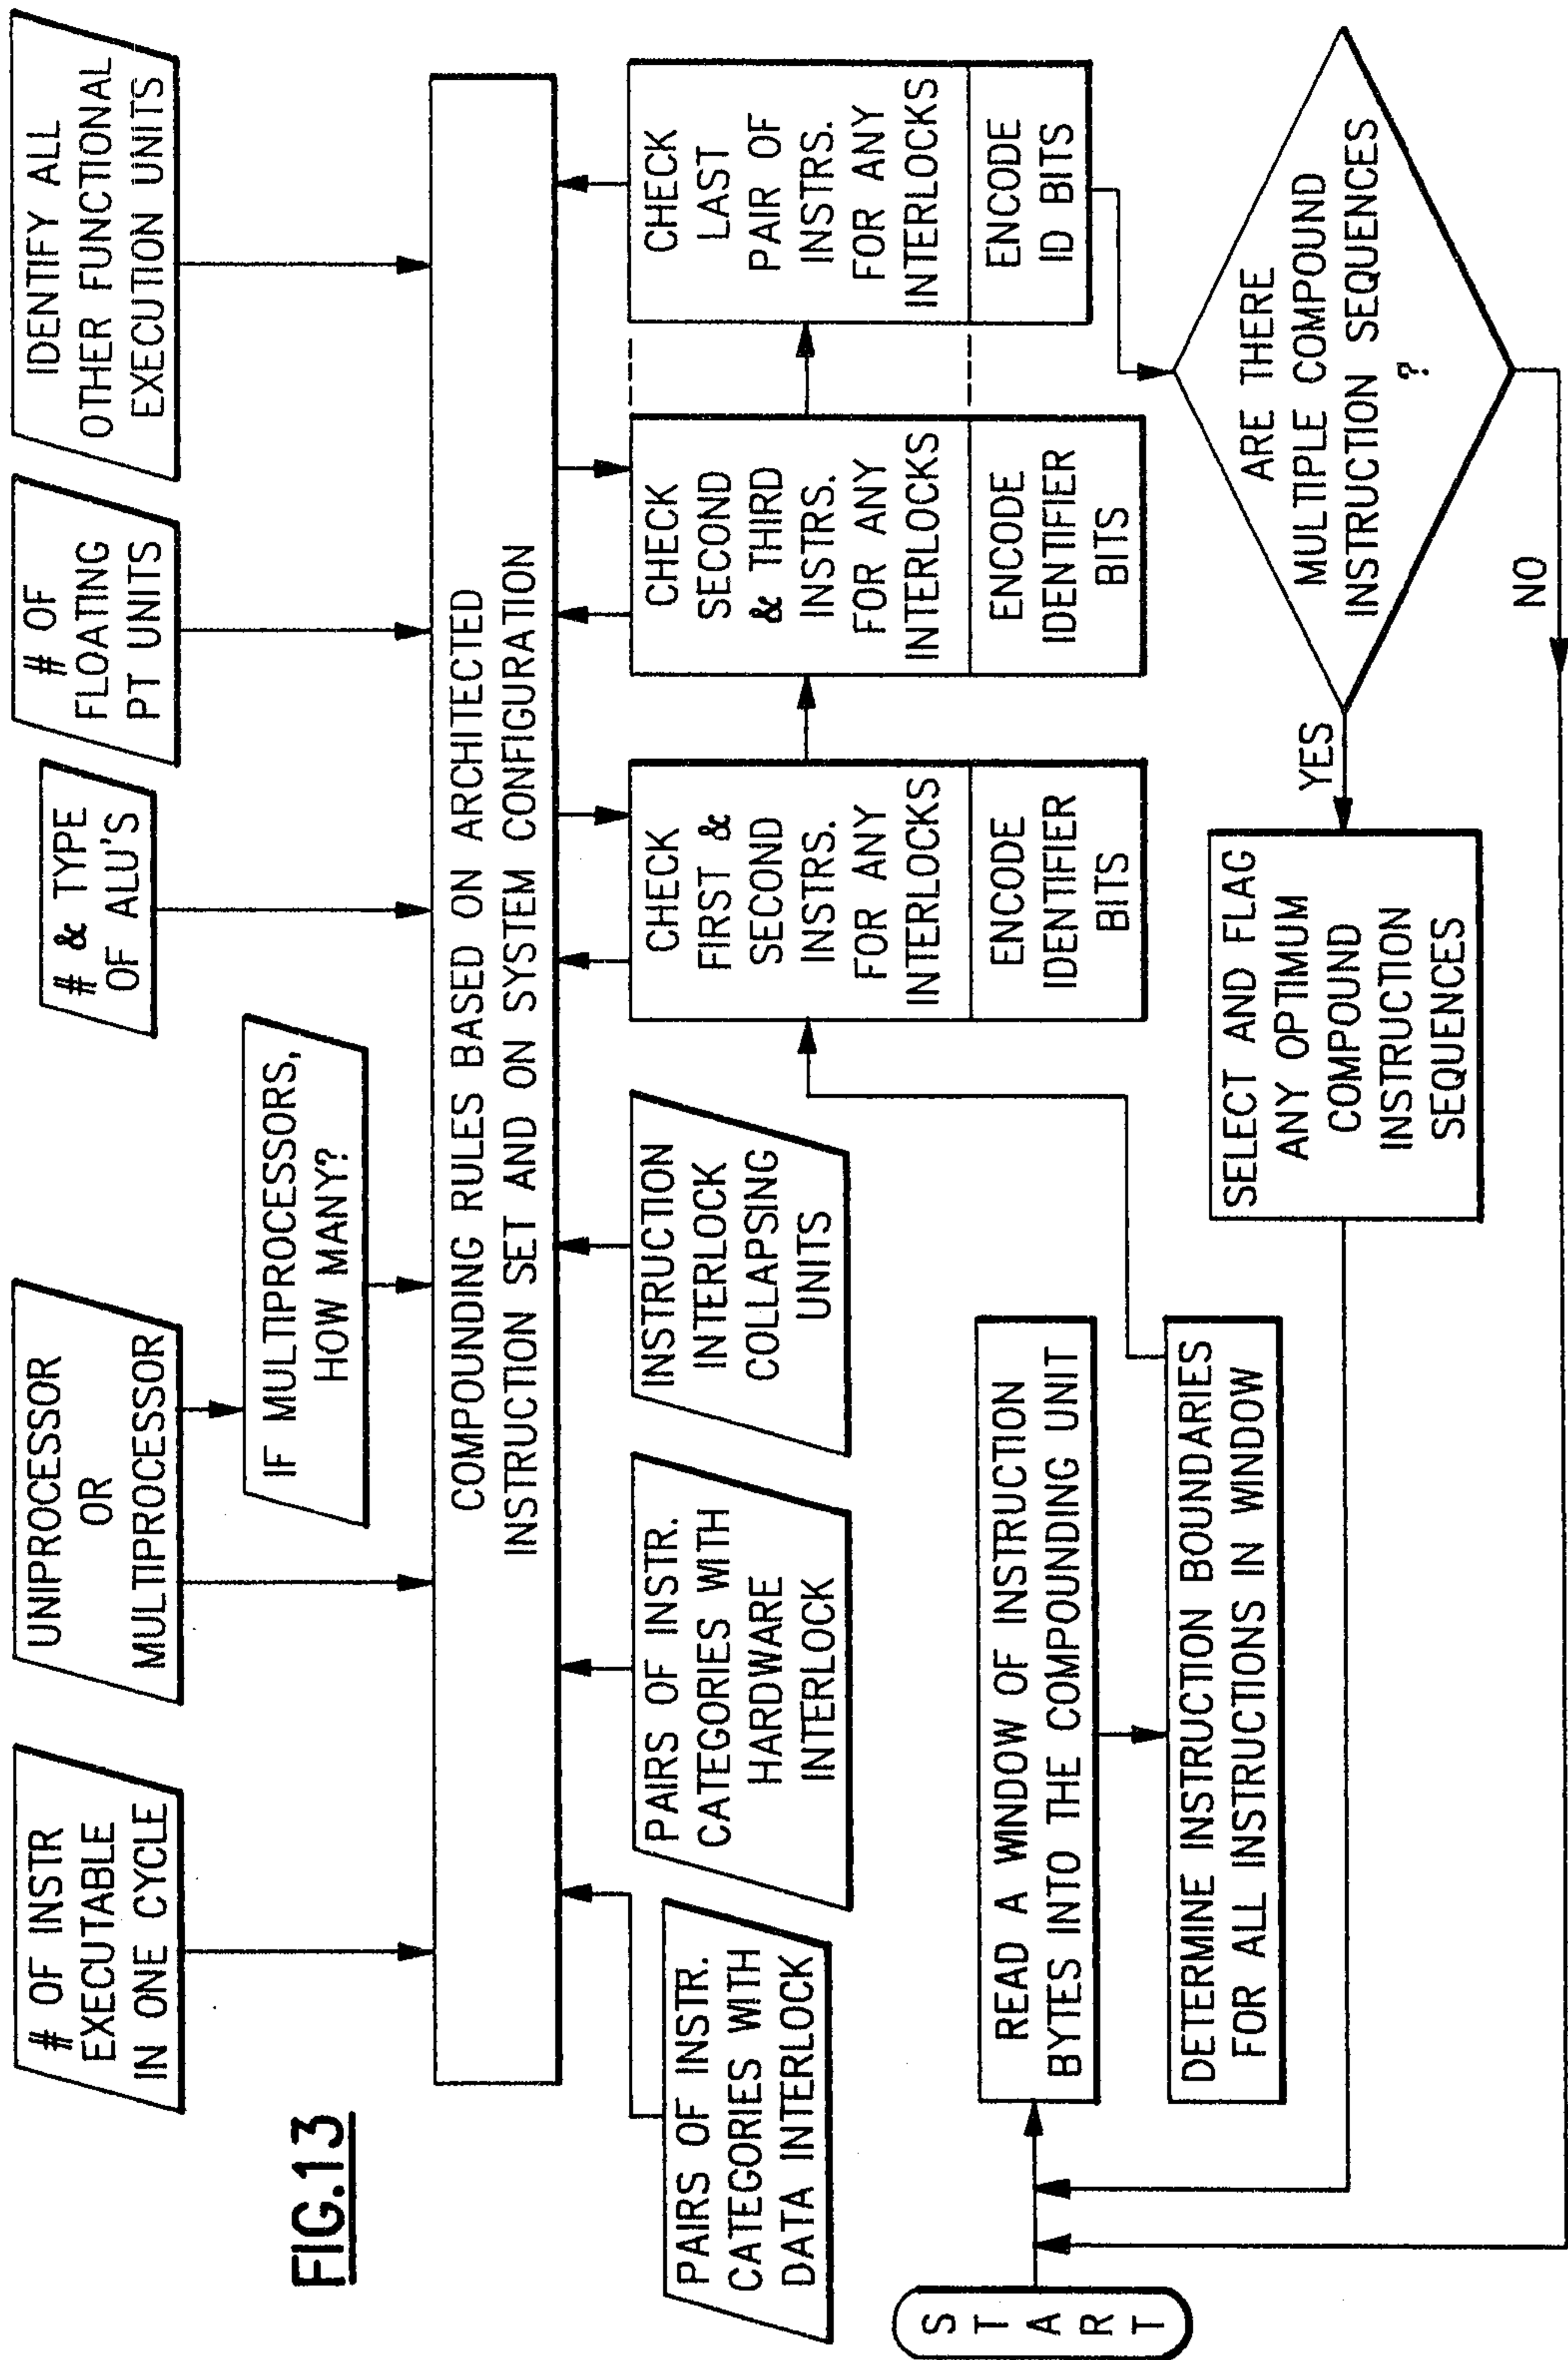

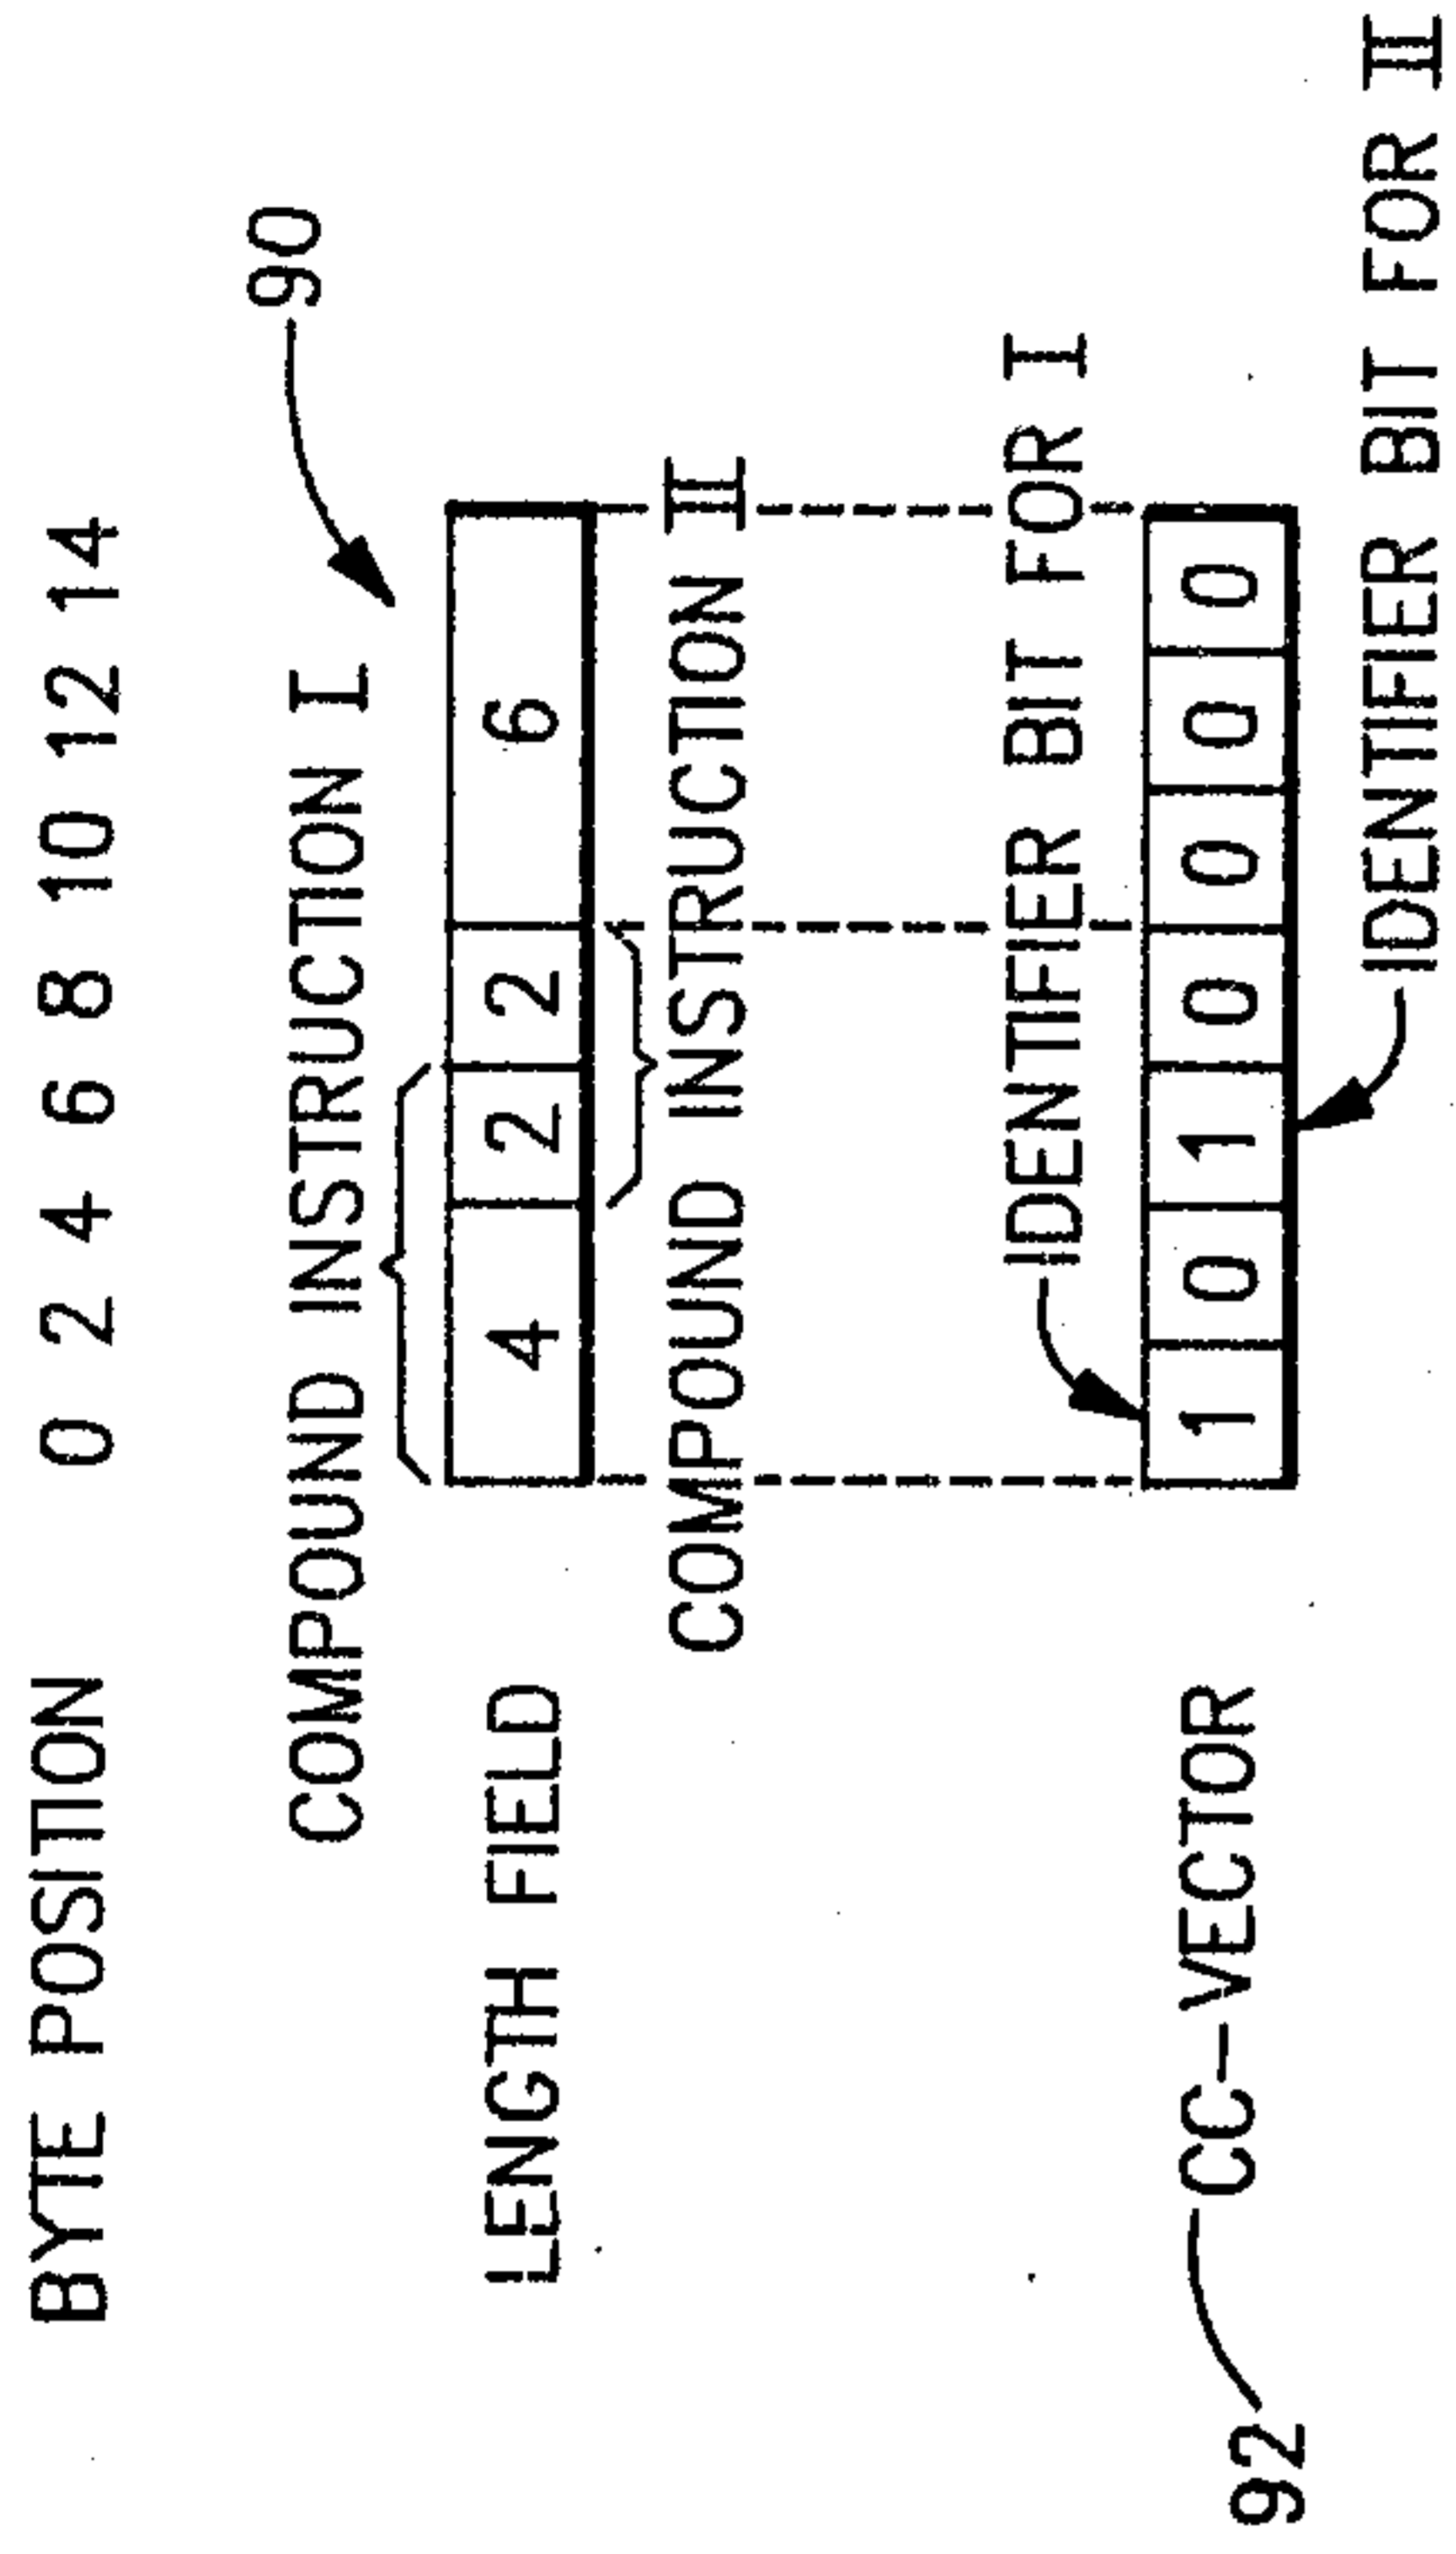


FIG.13

Alvin J. Fox
 PATENT AGENT

FIG. 14



Charles E. ...
PATENT AGENT




FIG.15

Alvin J. Jones
PATENT AGENT

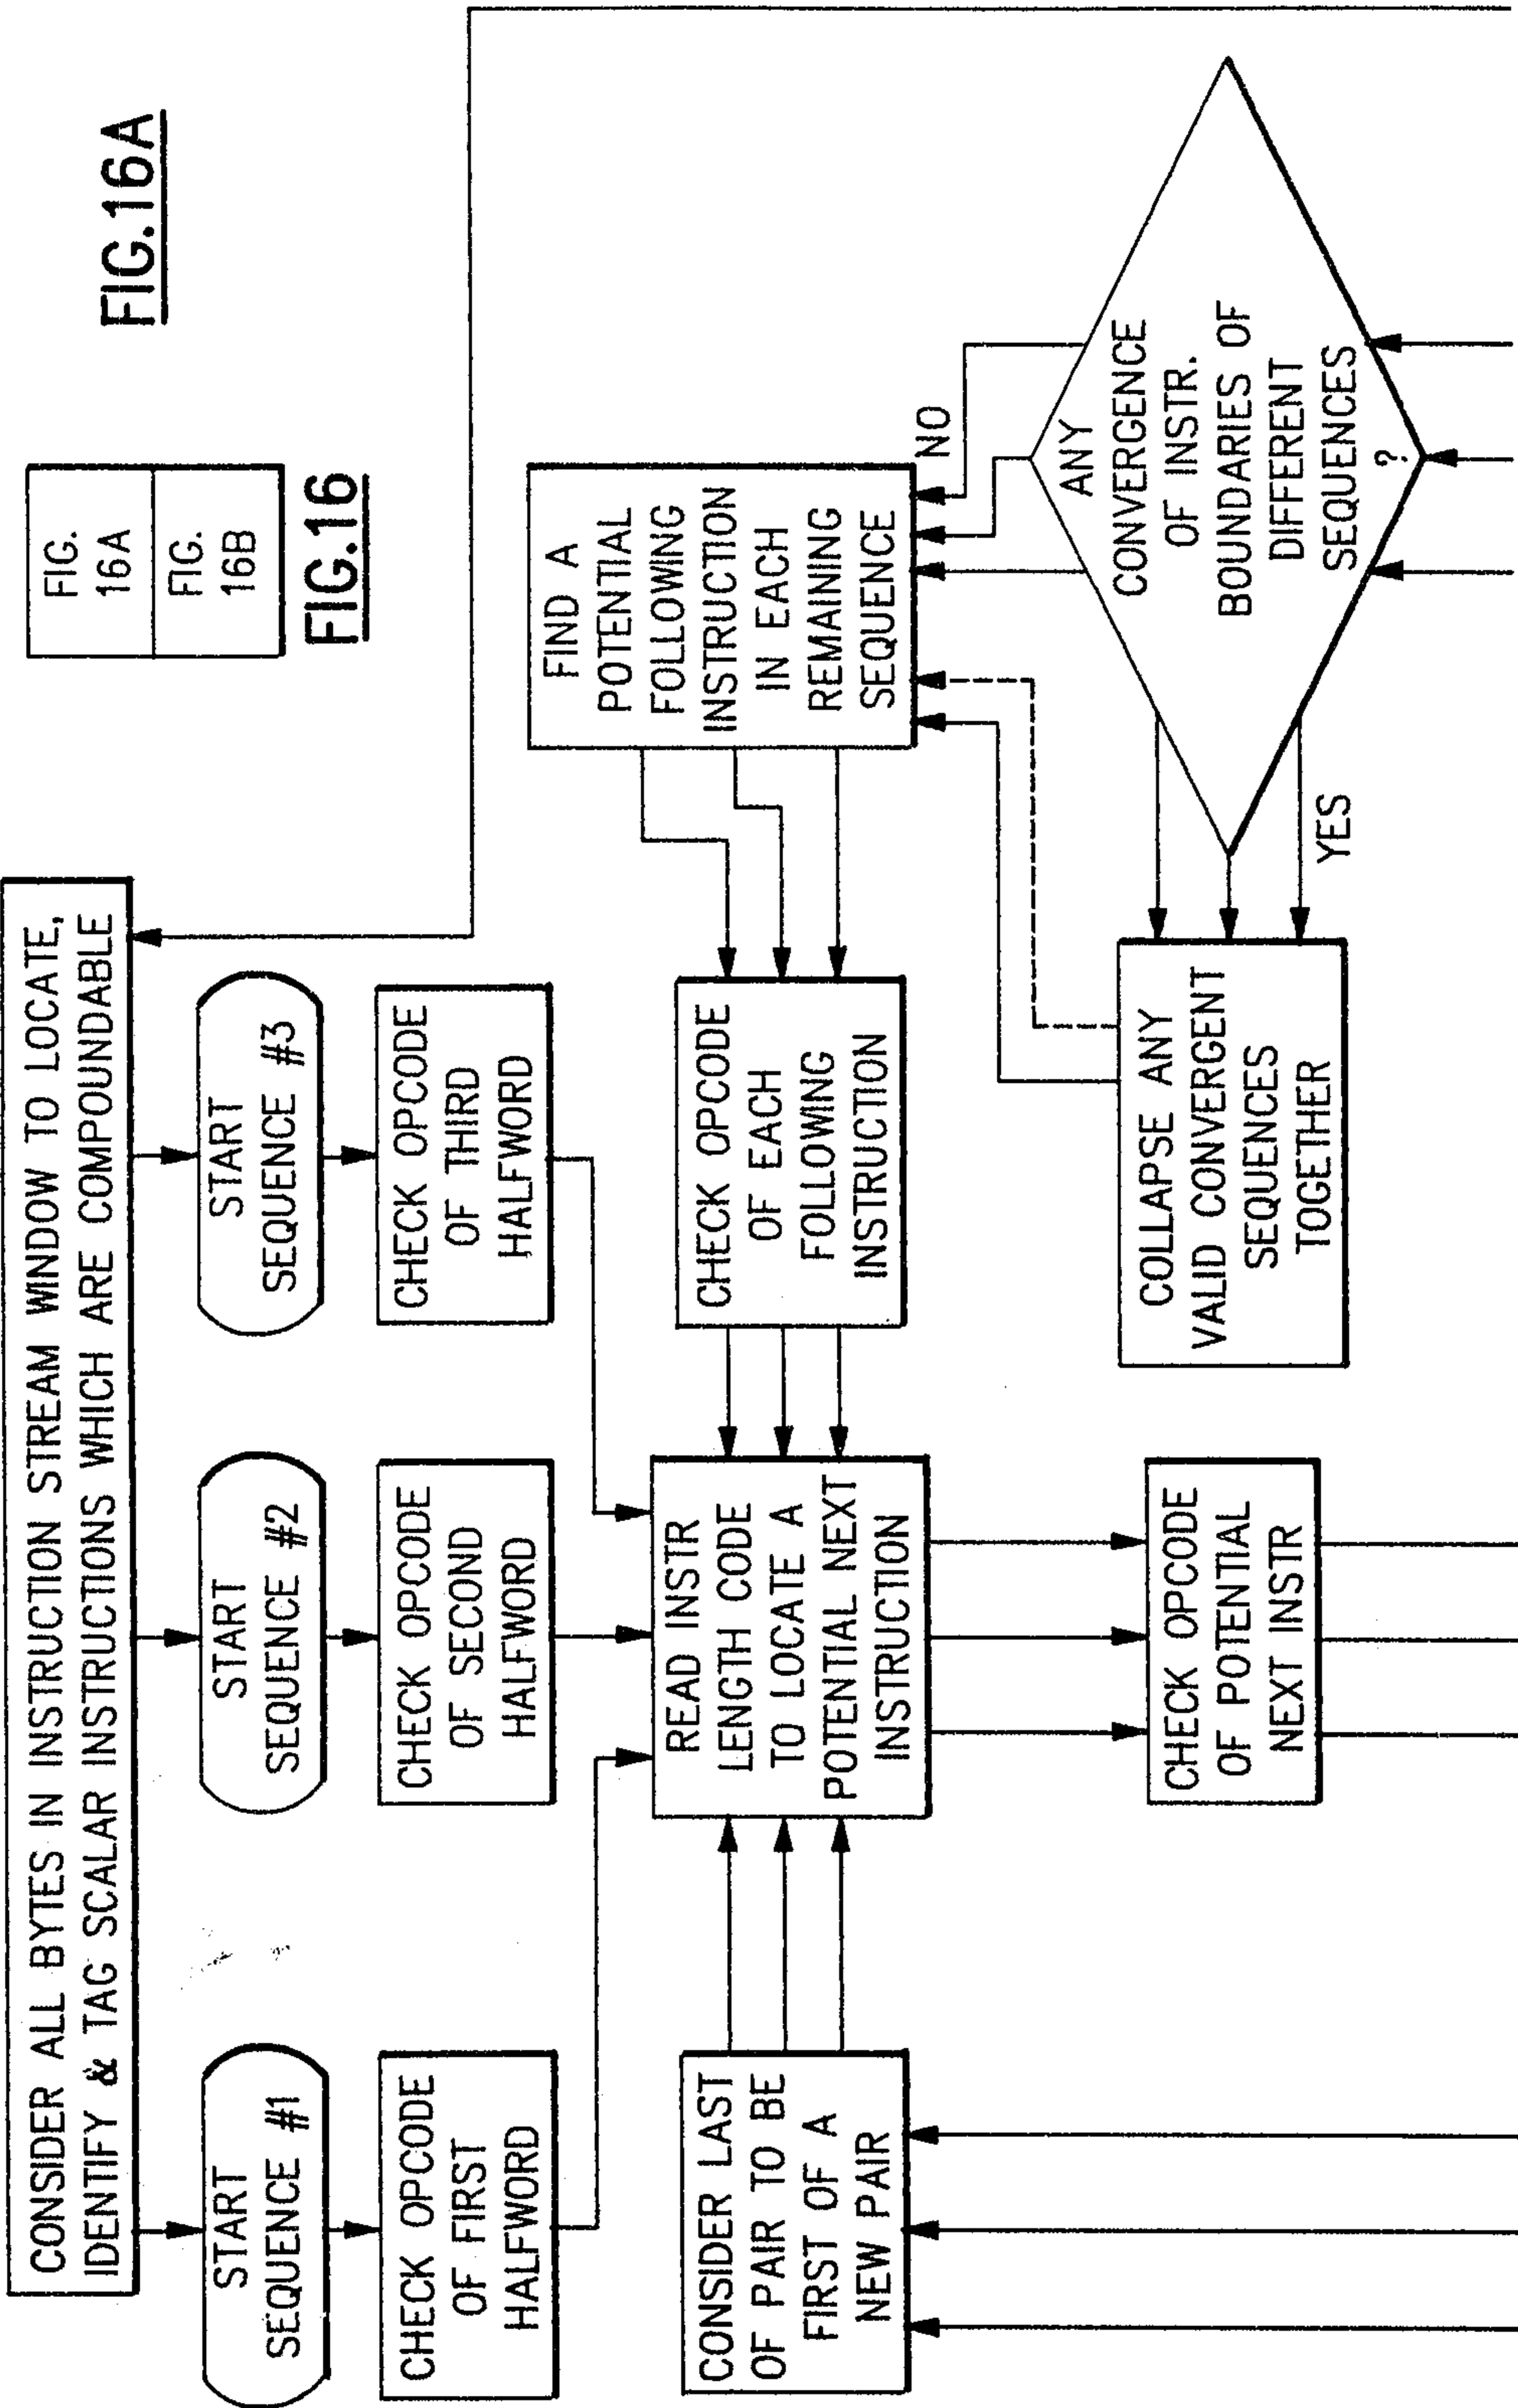


FIG. 16A

FIG. 16B

FIG. 16

[Signature]
 PATENT AGENT




FIG. 16B

Alvin J. ...
PATENT AGENT

CATE-GORY	1	2	3	4	5	6	7	8	9	10
1	A	S	A	A	A	S	A	A	A	A
2	S	N	S	S	S	N	S	S	S	S
3	A	A	N	A	A	A	A	A	A	A
4	A	A	A	A	A	A	A	A	A	A
5	A	A	A	A	A	A	A	A	A	A
6	A	N	A	A	A	N	N	A	N	N
7	S	S	S	S	S	N	N	S	N	N
8	A	A	S	S	S	S	S	S	S	A
9	A	S	S	S	S	S	S	A	N	N
10	A	A	A	A	A	A	A	A	N	N

KEY:

- A ALWAYS COMPOUND
- S SOMETIMES COMPOUND
- N NEVER COMPOUND

FIG.17

[Signature]
 PATENT AGENT