

(21) Application No: 0410986.4

(22) Date of Filing: 17.05.2004

(71) Applicant(s):  
**ARM Limited**  
(Incorporated in the United Kingdom)  
110 Fulbourn Road, Cherry Hinton,  
CAMBRIDGE, CB1 9NJ, United Kingdom

(72) Inventor(s):  
**Jan Guffens**  
**Ludwig Callewaert**  
**Koenraad Van Nieuwenhove**

(74) Agent and/or Address for Service:  
**D Young & Co**  
120 Holborn, LONDON, EC1N 2DY,  
United Kingdom

(51) INT CL<sup>7</sup>:  
**G06F 9/38**

(52) UK CL (Edition X ):  
**G4A AJR APX**

(56) Documents Cited:  
**US 5057837 A**

(58) Field of Search:  
UK CL (Edition W ) **G4A**  
INT CL<sup>7</sup> **G06F**  
Other: **Online:WPI,EPODOC,PAJ/JAPIO,TDB, INSPEC, XPESP**

(54) Abstract Title: **Program instruction compression**

(57) A processor is described including a plurality of data path elements 2, 4, 6, 8 (Fig.1) which independently perform in parallel different data processing operations. Program instructions are provided which are decoded to generate control signals for controlling the data path elements. Multiple instruction sets are supported with the same data processing operation to be performed by the same data path element being differently encoded within different instructions of different instruction sets. This enables code compaction when little parallelism may be achieved and full parallelism to be specified when this is possible.

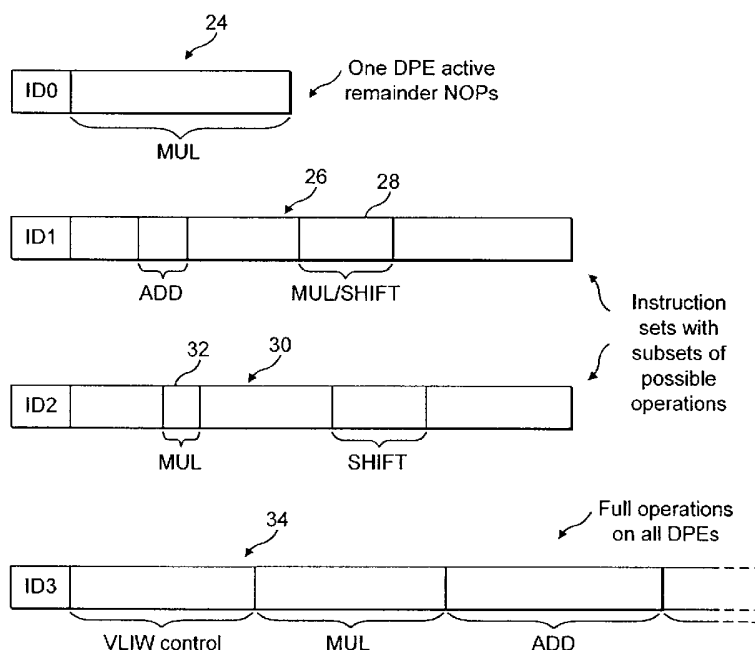


FIG. 4

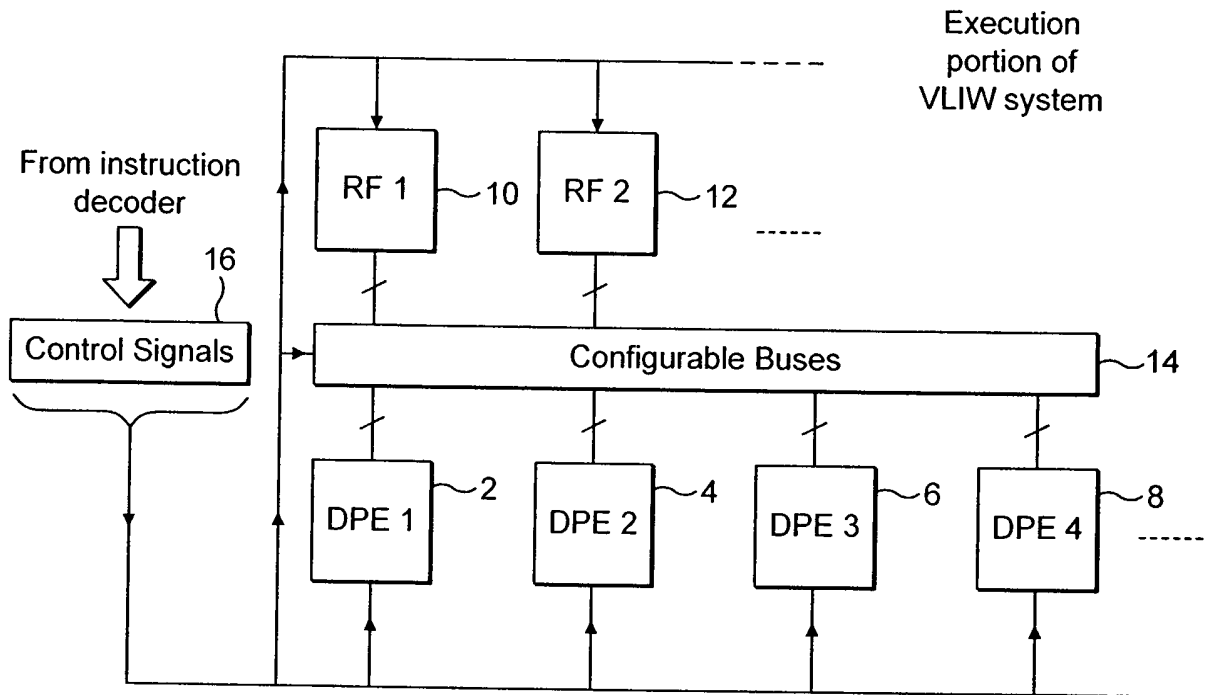


FIG. 1




FIG. 2

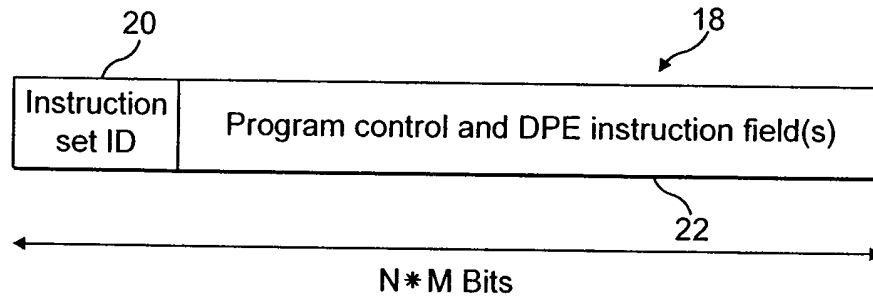


FIG. 3


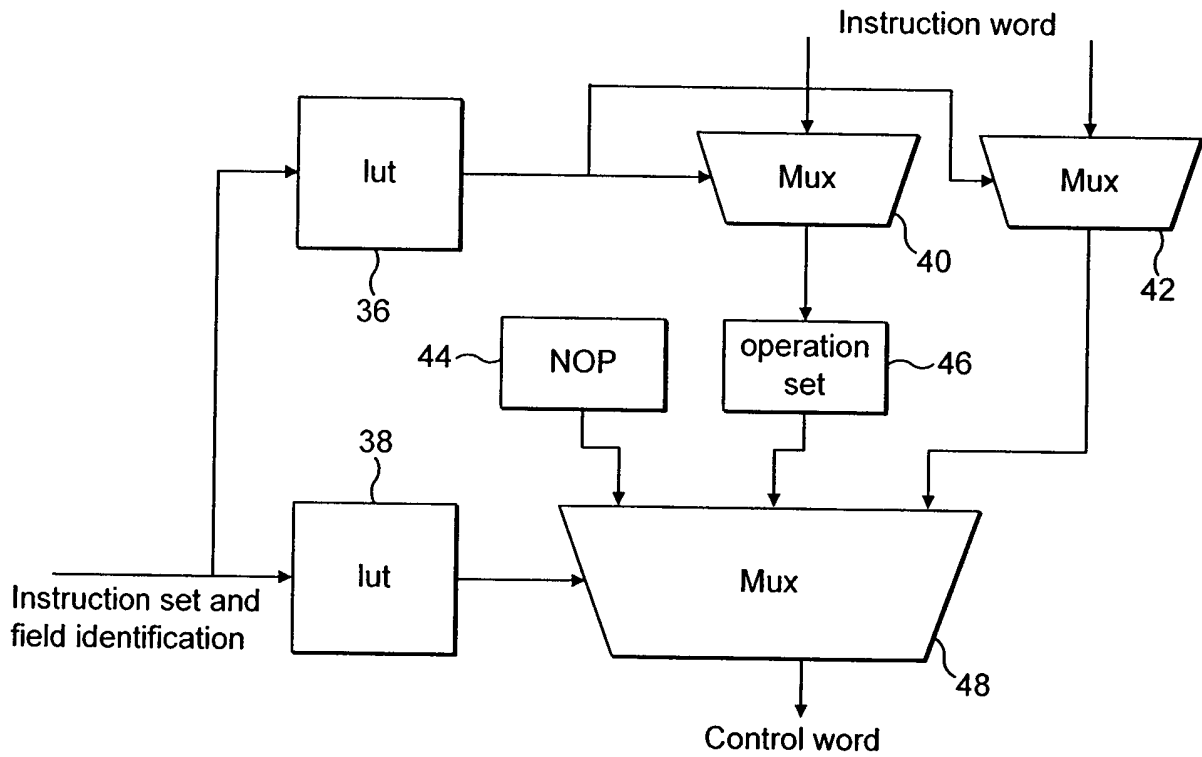



FIG. 4






FIG. 7

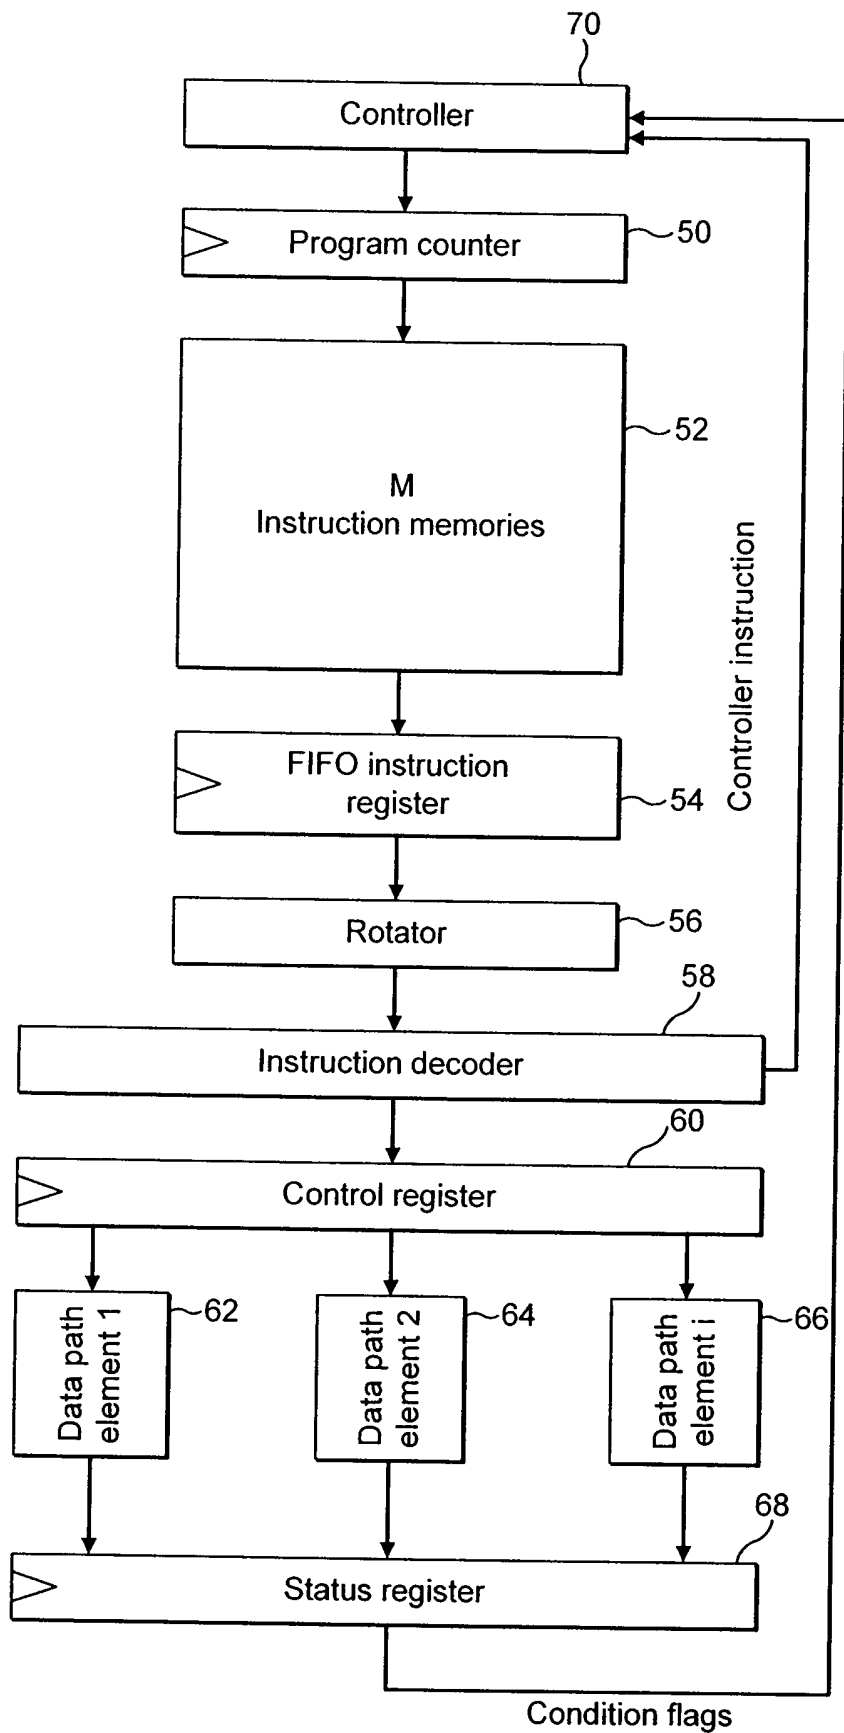


FIG. 8



**PROGRAM INSTRUCTION COMPRESSION**

This invention relates to the field of data processing. More particularly, this invention relates to data processing systems having a plurality of data path elements operable independently to perform in parallel respective data processing operations specified by a program instruction, such as, for example, so called very long instruction word (VLIW) processors and measures to reduce program code size for such systems..

The known TMS3206xx processor produced by Texas Instruments is designed for high-speed operation (e.g. 1 GHz) and consequently contains a simple instruction decoder. This processor uses 32-bit instructions. Instructions are loaded from a memory in a 256-bit fetch packet containing eight 32-bit instructions. Each instruction contains a bit (the P bit) that indicates if the next instruction in the fetch packet can be executed in the same clock cycle. Instructions that execute in the same clock cycle are called the execute packet. Since an execute packet cannot cross a fetch packet boundary, the P bit of the last instruction in the fetch packet must be cleared. If a functional unit within the processor is not addressed by an instruction within the execute packet, then it performs a default operation, such as a Nop.

The SC140 Processor produced by StarCore builds its instruction words up out of 16-bit words. Most instructions consist of a single 16-bit instruction word. Some instructions need two instruction words. An instruction prefix word (16 or 32 bits) can be specified. This prefix is used to extend the number of addressable register fields, conditionally execute instructions (guarded execution), or to specify the number of instructions to be executed in one clock cycle. If no prefix word is used, then the instructions are linked together using a bit in the instructions in a similar way to the TMS320C6xx processor discussed above.

Within the SC140 processor, instructions are fetched from the memory in 128-bit units (8\*16-bit words). Up to six functional units can be controlled in one clock



cycle. The instructions that execute in one clock cycle can span a 128-bit boundary. An instruction alignment circuit performs necessary alignment operations when the instructions span such a boundary.

5           The Thumb enabled scalar processors produced by ARM Limited are able to execute either 32-bit ARM code or 16-bit Thumb code. The Thumb instruction set does not provide all instructions that can be specified within the ARM instruction set.

10           VLIW processors such as the TMS320C6xx and SC140 processors are advantageous in providing for highly parallel execution of data processing operations. However, as the complexity of processing operations to be performed steadily increases, the high program memory storage requirements associated with these VLIW processors become a significant disadvantage.

15           Viewed from one aspect the present invention provides apparatus for processing data, said apparatus comprising:

          an instruction decoder operable to decode one or more instruction fields within a program instruction to generate at least control signals; and

20           a plurality of data path elements responsive to said controls signals independently to perform in parallel respective data processing operations specified by said program instruction, at least some program instructions decoded by said instruction decoder allowing control signals for respective data path elements to be independently specified; wherein

          said instruction decoder is operable such that a given data processing operation  
25   to be performed by a data path element can be specified using differently encoded instruction fields in different program instructions.

          The present technique recognises that in many real life processing applications to be performed, the high levels of parallelism provided by the very long instructions supported by VLIW processors are not exploited. As an example, whilst high degrees  
30   of parallelism may be achieved within a tight program loop performing digital signal processing type operations upon a large number of independent data values, such as

pixel values, significant amounts of time will be spent executing control code and other code types in which there is little potential for parallelism. Having recognised these differing types of real life operation, the present technique provides multiple instruction encodings whereby the different instruction encodings that specify the same data processing operation can be targeted at different operational regimes. As an example, within control flow types of operation, it may be apparent that it is very common for only a single data path element to be active at any given time and accordingly smaller compressed program instructions can be used which only allow a single data path element to be active. Considerable program memory space can be saved by the assumption that all the other data path elements will take some specified default behaviour when a program instruction does not indicate a data processing operation which they are to perform. Conversely, when a portion of a computer program is encountered in which high degrees of parallelism may be exploited, then a different program instruction encoding may be used to enable a wider range of the data path elements to be simultaneously controlled and a wider range of data processing operations upon each of those data path elements to be specified. It is also possible that some intermediate length encodings will provide for the commonly used operations (e.g. add) of certain datapath elements, but not the less commonly used operations (e.g. bitwise invert).

Whilst it will be appreciated that the differently encoded respective instruction fields for the independent data path elements could vary their encoding in a variety of different manners, particularly useful ways of varying the encoding to improve code compression are to vary the bit lengths used for the instruction fields and the bit positions within the program instructions of the instruction fields.

In preferred embodiments the program instructions are divided into a plurality of instruction sets and the instruction decoder is responsive to an instruction set identifying field within the program instructions to determine which bits within the program instruction correspond to respective instruction fields for different data path elements. Accordingly, a plurality of different instruction sets may be provided suited to different operational regimes, such as a full functionality instruction set with large instructions suitable for highly parallel operation, a highly compressed instruction set

in which only one data path element may be operative at any one time suited to highly non-parallel operation and intermediate instruction sets in which forms parallelism commonly encountered is supported but not the full level of potential parallelism.

5 The instruction set decoder is advantageously responsive to the instruction set identifier to determine program instruction length in order to deal with variable program instruction lengths within the system.

The differently encoded instruction fields serve to represent different sets of data processing operations that can be performed by the respective data path elements.  
10 Desirably some of these sets are subsets of others and include the most commonly occurring processing operations which it is desired to specify.

Another preferred way of improving program instruction compression is to support reduced operand ranges within the smaller instructions, such as a reduced  
15 range of register specifiers, a reduced range of jump addresses for program branches, a reduced set of functionality (e.g. fewer supported operation types) or a reduced range of immediate values.

An instruction field may relate to only a single possible data path element.  
20 Alternatively, improved compression can be achieved in some circumstances when an instruction set identifier itself specifies which data path element it controls. This effectively yields a form of tree encoding/decoding structure which may be readily adapted to the particular circumstances of the processors operation.

25 It will be appreciated that with different program instructions, differing numbers of data path elements will be active and inactive. Data path elements which are not specified a particular processing operation by a given program instruction will perform a default operation, such as a Nop, or some pending operation, such as a multicycle operation that was previously initiated.

30

The efficiency with which the program may be stored within memory is improved in embodiments in which the program instructions are readable in fixed

length memory accesses which contain at least portions of a plurality of program instructions and a rotating and address incrementing mechanism is used to align program instructions for supply to the instruction decoder. The degree of alignment required can be determined from the instruction set identifiers which themselves specify or imply an instruction length. Furthermore, it is possible to split the memory into different banks with the program spread across those banks such that only those banks which require reading to replace portions of program instructions that have been consumed need be powered thus saving power.

10 Viewed from another aspect the present invention provides a method of processing data, said method comprising the steps of:

decoding one or more instruction fields within a program instruction to generate at least control signals; and

15 in response to said controls signals controlling a plurality of data path elements independently to perform in parallel respective data processing operations specified by said program instruction, at least some program instructions allowing control signals for respective data path elements to be independently specified; wherein

20 a given data processing operation to be performed by a data path element can be specified using differently encoded instruction fields in different program instructions.

Viewed from a further aspect the present invention provides a computer program product for controlling an apparatus for processing data having an instruction decoder operable to decode one or more instruction fields within a program instruction to generate at least control signals and a plurality of data path elements responsive to said controls signals independently to perform in parallel respective data processing operations specified by said program instruction; said computer program product comprising:

30 at least one program instruction including an instruction field with a first encoding specifying a given data processing operation to be performed by a data path element; and

at least one different program instruction including a differently encoded instruction field with a second encoding different to said first encoding and specifying said given data processing operation to be performed by said data path element.

5 Viewed from a complementary aspect the present technique may also be considered to be a compiler program for producing computer programs in accordance with the above described features. The compiler effectively has a plurality of instruction sets from among which it may select program instructions to perform the data processing operations for which it is attempting to compile code.

10

Embodiments of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1 schematically illustrates the data path of a VLIW processing system;

15 Figure 2 schematically illustrates the degree of parallelism which may be achieved at different instruction cycles within data processing system operation;

Figure 3 schematically illustrates a generic program instruction for controlling a VLIW processor that is suitable for use with the current techniques;

20 Figure 4 illustrates four example program instructions having four respective differently encoded instruction fields used to specify a given data processing operation to be performed by a given data path element;

Figure 5 schematically illustrates a decoding/encoding tree which may be used as part of an instruction compression technique;

Figure 6 schematically illustrates a portion for one control word of an instruction decoder;

25 Figure 7 schematically illustrates one example of the layout of program instructions within a memory which is read with fixed width memory accesses;

Figure 8 schematically illustrates a VLIW processor supporting multiple instruction sets, some of which are decoded to decompress them; and

30 Figure 9 schematically illustrates the circuit for fetching and aligning program instructions from a memory.

Figure 1 shows a portion of a VLIW processor including a plurality of independently controllable data path elements 2, 4, 6, 8, which may, for example, be elements such as an adder, a multiplier, a shifter, an alu, a floating point unit, or other functional unit. These data path elements 2, 4, 6, 8, are connected to the register files 10, 12 via a controllable bus network 14. Control signals stored within a control signal register 16 configure the data path elements 2, 4, 6, 8, the register files 10, 12 and the controllable bus network 14 to perform a desired data processing operation or parallel set of operations in a given processor clock cycle. It will be appreciated that typically a large number of control signals are needed to control all these circuit elements and accordingly the control signal register 16 is comparatively wide. Simply storing the control signals within an instruction memory as a program whilst possible is highly inefficient in terms of program memory storage requirements. The control signals within the control signal register 16 are in Figure 1 derived from an instruction decoder. This instruction decoder, as will be described later, decompresses program instructions that represent desired behaviour in a compact form to yield the full set of control signals needed to control the circuit element illustrated in Figure 1.

Figure 2 is a diagram schematically illustrating different possible independent data processing operations which can be performed with different program instructions. As will be seen, in the first section of the instruction cycles of Figure 2 parallelism is not generally possible with each program instruction being dependent upon completion of a previous instruction, either through a data dependency or by virtue of requiring the same resource. In the middle section of Figure 2 a portion of code is encountered in which highly parallel operation is possible, such as the parallel processing of pixel or other independent signal values. Within this regime, the full flexibility and complexity of data processing operations which may be specified by full width VLIW instruction words is desirable and advantageous. In the final portion of Figure 2 the system again returns to a regime in which parallel operation is not exploitable.

Figure 3 schematically illustrates a generic program instruction according to one example of the present technique. This program instruction 18 includes an

instruction identifier field 20 and a section 22 which includes bits specifying control of the VLIW controller itself as well as instruction fields relating to one or more data path elements which it may be desired to use. In this example embodiment, the memory system which stores the program instructions is accessible in N-bit blocks with M such blocks being accessible in parallel. Accordingly, the maximum width of any program instruction 18 is  $M \times N$  bits. In practice it is convenient to constrain the program instructions to the multiples of N bits in length these all being less than or equal to  $M \times N$  bits in length.

Figure 4 schematically illustrates four different program instructions from respective different instruction sets which each include an instruction field that can specify a given data path element operation to be performed using different encodings. In program instruction 24, the instruction set is one in which only a single data path element may be addressed by a single program instruction. The program instruction includes within its single instruction field a specifier of which data path element is being addressed as well as the necessary opcodes and operands to control that data path element. The remaining data path elements are either inactive and perform Nops or perform some other default operation, such as continuing a previously started multi-cycle processing operation or responding to some previously established state.

The second program instruction 26 is longer and allows multiple data path elements to be addressed within a single instruction. The bit position and bit length of the instruction field 28 which relates to the multiplier varies compared to that in program instruction 24 and the other instruction. The instruction field length is less and accordingly only a subset of data processing operations of which the multiplier is capable may be specified within this shortened instruction field. The register specifier range may be reduced and/or the immediate value range may be reduced in order to facilitate this shortening of the instruction field length. The more commonly required instructions are chosen to be supported with the subset. In other example instructions, a particularly preferred type of variable for restriction to enable better compression is the jump specifier range since full address space jumps are rare and it is much more

common to have relatively short jumps within the program instruction flow. Within this second instruction 26, the instruction field 28 is not dedicated to the multiplier and includes bits which specify whether either the multiplier or the shifter is being addressed by the instruction field concerned.

5

In the third program instruction 30, a dedicated instruction field 32 is provided which only addresses the multiplier. This dedicated instruction field 32 is smaller still than the instruction field 28 and is at a different bit position.

10 Finally the fourth illustrated program instruction 34 fully addresses all of the data path elements with individual instruction fields as well as including a VLIW control field which specifies control parameters for the VLIW controller itself.

Figure 5 is a diagram schematically illustrating a decoding tree that may be used to decode an input instruction field within a program instruction. At each node the decoding branches to either a definitely specified data path element which is being  
15 addressed or a further node at which some further branching can take place. Ultimately, as the decoding tree is traversed the branches will terminate at a specified data path element.

20 Figure 6 schematically illustrates an instruction decoder slice for one data path element within the processor which serves to decode/decompress the program instructions to form the required control signals. The instruction set identifier and field identification is supplied to two respective look-up-tables 36, 38 which in turn generate multiplexer controlling signals. The block within the decoder comprising the  
25 multiplexer 42 serves as an instruction extractor and extracts an appropriate instruction field from within the program instruction concerned. It will be appreciated that the bit position of the instruction field for a given data path element varies for the different encodings and the different instruction sets and accordingly the multiplexers 40, 42 select from across the full span of the program instruction word to pick out the  
30 desired instruction field.



The second portion of the instruction decoder takes the selected instruction field and expands it as required to form the full set of control signals required for that data path element. It may be that the multiplexer 42 directly provides this full set of control signals, or alternatively if default Nop operation is required then these signals may be generated by the circuit element 44. In the case of an instruction field which is representing a subset of possible data path element operations using a compressed encoding of the control signals, then the circuit element 46 can expand this compressed representation to form the full set of control signals which are then selected by the multiplexer 48.

Figure 7 schematically illustrates a memory for storing program instructions implemented as a set of M memories with a width of N bits each. These M memories can all be read in parallel. Between one and M words are read from the memories for every program step. The words that are read are stored in the instruction register 54 (illustrated in Figure 9). The instruction register acts as a FIFO for code words. The instruction extraction then consists of a read of program words from the instruction register and using a rotator to align these for instruction decoding. Figure 7 is an example of the distribution of program words over the different memories.

The decoding flow for the instructions shown in Figure 7 is:

1. cycle 1 :
  - a. Read 3 code words of instruction 1 and 1 word of instruction 2 into the instruction register.
2. cycle 2 :
  - a. Read 2 code words of instruction 2 and 1 code word of instruction 3 into the instruction register.
  - b. Use the rotator to MSB align instruction 1 at the input of the instruction decoder by rotating it over 1 word and decode it (this is an example with one endianness, but the technique is not limited to such systems).
3. cycle 3 :
  - a. Read 2 additional code words of instruction 3 and one code word of instruction 4 into the instruction register.

- b. Use the rotator to MSB align instruction 2 at the input of the instruction decoder by rotating it over 2 words and decode it.
- 4. cycle 4 :
  - a. Read 2 code words of instruction 4 and instruction 5 into the instruction register.
  - b. Use the rotator to MSB align instruction 3 at the input of the instruction decoder by rotating it over 3 words and decode it.
- 5. cycle 5 :
  - a. Read 3 code words of instruction 6 into the instruction register.
  - b. Decode instruction 4.
- 6. cycle 5 :
  - a. Read the last code word of instruction 6 into the instruction register.
  - b. Use the rotator to MSB align instruction 5 at the input of the instruction decoder by rotating it over 3 words and decode it.
- 7. ...

Figure 8 is a schematic high level diagram illustrating an example embodiment of the VLIW processor using the present techniques. A program counter 50 is used to address an instruction memory 52 storing the variable length program instructions. These variable length program instructions are read out as fixed length memory access words into a FIFO instruction register 54 from which the individual program instructions are extracted and aligned by a rotator 56 before being supplied to an instruction decoder 58. The instruction decoder 58 decompresses, as appropriate, these program instructions to produce control signals stored within a control register 60. The control signals configure the different data path elements 62, 64, 66 to perform a desired processing operation with the results being used to update a status register 68 which in turn feeds back to a VLIW controller 70. The VLIW controller 70 controls the program counter to increment appropriately through the program stored within the memory 52.

Figure 9 schematically illustrates one example structure of the controller 70. It will be seen that both the current program counter value and an incremented version

of this are maintained such that appropriate memory addresses to address the different memories may be generated when a desired access is unaligned and wrapped around the different memories. Instruction length derived by the instruction set decoder 58 is used to control the program counter advance as well as any required jump commands forcing a jumping program flow control.

The following is an example illustrative of the possible subsets of Alu data path element controls which it may be desired to support within different instruction sets. Other data path elements may have their instruction sets selected/configured in a similar way.

**- This Alu has 4 control ports that can be independently controlled: cmd0, cmd1, cmd2 and cmd3.** When only parts of the ALU are used in an application it doesn't make sense driving all of its parts from the microcode. The user can use instruction set definitions for this purpose.

The alu consists of multiple independent units controlled by a separate commando busses. For each of these commando busses independent an operation set must be defined. The 'full' instruction set of an ALU is generated for ALU "alu\_1" as follows when the user has not specified any instruction sets:

```
operation_set("full","alu_1:cmd0",["add","sub","passx","passy",
                                     "incx","incy","decx","decy","aslx",
                                     "asly","rsub","negx","negy","exor",
                                     "exnor","or","and","absx","absy",
                                     "min","max","comx","comy","varshiftright",
                                     "varshiftright","const","default"]); // 5
                                     bits
30 operation_set("full","alu_1:cmd1",["zero","one","two","three",
                                     "four","five","six","seven","eight",
                                     "nine","ten","eleven","twelve","thirteen",
```

```

“fourteen”, “fifteen”, “thirtyone”, “sixtythree”, “fgt”,
“flt”, “fge”, “fle”, “feq”, “fne”,
“fgtz”, “fltzt”, “fgez”, “flez”, “feqz”,
“fnez”, “rslt”, “<#----->”, “default”]); // 6

```

5 bits

```

operation_set(“full”, “alu_1:cmd2”, [“upsh”, “downsh”, “pass”, “immediate”]); // 2 bits
operation_set(“full”, “alu_1:cmd3”, [“signed”, “unsigned”, “X”, “Y”, “default”]); // 1 bits

```

10 The code width of an operation set is determined by the number of instruction in the set. E.g.  $\text{ceil}(\log_2(\# \text{instructions in set}))$ .

```

control_group(“CG_alu_1”, [“reg_dx_alu_1:ar0:full”,
                           “reg_dy_alu_1:ar0:full”,
                           “alu_1:cmd0:full”,
15      “alu_1:cmd1:full”,
                           “alu_1:cmd2:full”,
                           “alu_1:cmd3:full”]);

```

20 So, the full instruction set of an ALU is represented by a single control\_group “CG\_alu\_1” which consists of the read addresses of the registers connected to the input pins of the ALU, and of the ‘full’ operation sets of all control ports of the ALU.

The width of a control group is the sum of the widths of the contained operation sets.

25 In order to be able to define instruction sets with restricted functionality on the ALU new operation sets and control groups need to be defined. Before this can be done, the user needs to determine which part of the functionality he wants in a specific instruction set.

30 **Example 1:** When the ALU is only used as a flag control unit, the following definitions should be used:

```

operation_set("fcu", "alu_1:cmd0", ["passx", "passy", "exor", "exnor",
                                     "or", "and", "comx", "comy"]); // 3 bits
operation_set("default", "alu_1:cmd1", ["default"]); // 0 bits
operation_set("pass", "alu_1:cmd2", ["pass"]); // 0 bits
5 operation_set("default", "alu_1:cmd3", ["default"]); // 0 bits

control_group("CG_fcu_alu_1", ["reg_dx_alu_1:ar0:full",
                              "reg_dy_alu_1:ar0:full",
                              "alu_1:cmd0:fcu",
10 "alu_1:cmd1:default",
                              "alu_1:cmd2:pass",
                              "alu_1:cmd3:default"]);

```

Control group "CG\_fcu\_alu\_1" will only contain the read addresses of the input  
15 registers and 3 bits to store the instructions for the 'cmd0' control port. The  
instruction set decoder will automatically apply the default instructions to the 'cmd1  
and the cmd3 control ports and the pass instruction to the 'cmd2 control port.

**Example 2:** The instruction set for the ALU without the shift unit is specified as  
20 follows:

```

operation_set("noshift", "alu_1:cmd0", ["add", "sub", "passx", "passy",
                                         "incx", "incy", "decx", "decy", "aslx",
                                         "asly", "rsub", "negx", "negy", "exor",
25 "exnor", "or", "and", "absx", "absy",
                                         "min", "max", "comx", "comy", "default"]); // 5 bits

operation_set("flags", "alu_1:cmd1", ["fgt", "flt", "fge", "fle", "feq", "fne",
                                       "fgtz", "fltz", "fgez", "flez", "feqz",
30 "fnez", "rslt", "default"]); // 4 bits

control_group("CG_alu_noshift", ["reg_dx_alu_1:ar0:full " ,

```

```

“reg_dy_alu_1:ar0:full”,
“alu_1:cmd0:noshift”,           // 5 bits
“alu_1:cmd1:flags”,           // 4 bits
“alu_1:cmd2:pass”,           // 0 bits
5 “alu_1:cmd3:full”]);        // 1 bit

```

A control bit can be gained by defining a new operation set on the ‘cmd’ control port.  
For example, for logical =0 and absminmax =0:

```

10 operation_set(“logical0absminmax0”,”alu_1:cmd0”,[“add”,”sub”,”passx”,
“passy”, “incx”,”incy”,”decx”,”decy”,”aslx”,
“asly”,”rsub”,”negx”,”negy”,”default”]); // 4
bits

```

15 Operation set “alu\_1:cmd:full” can then be replaced in above control\_group definition  
by “alu\_1:cmd0:logical0absminmax0”.

**Example 3:** The ALU\_NOFLAG can be described by the following definitions.

```

20 control_group(“CG_alu_noflag”,[ “reg_dx_alu_1:ar0:full”,
“reg_dy_alu_1:ar0:full”,
“alu_1:cmd0:noshift”,           // 5 bits
“alu_1:cmd1:default”,           // 0 bits
25 “alu_1:cmd2:pass”,           // 0 bits
“alu_1:cmd3:full”]);          // 1 bit

```

30 The following example illustrates the definition of a complete instruction set decoder  
for a processor with three datapath elements : 1 ALU, 1 RAM and an constant  
generator. The following operation sets are defined:

```

operation_set("full","ram_1:rwn",["write","read","default"]); // 2 bits
operation_set("full","immediate_1:immediate",["<#----->"]); // 16 bits
operation_set("full","alu_1:cmd0",["add","sub","passx","passy",
5      "incx","incy","decx","decy",
      "rsub","negx","negy","exor",
      "exnor","or","and","default"]); // 4 bits
operation_set("addsub","alu_1:cmd0",["add","sub","incx","decx","default"]); // 2 bits

```

Control groups are simple groupings of control words that will be dealt with in identical ways in the instruction set definitions. The following control groups are defined :

```

control_group("CG_immediate_1",["immediate_1:immediate:full"]); // 16 bits
control_group("CG_ram_1",["ram_1:rwn:full"]); // 2 bits
15 control_group("CG_alu_addsub",["alu_1:cmd0:addsub"]); // 2 bits
control_group("CG_alu_full_1",["alu_1:cmd0:full"]); // 4 bits

```

Instruction fields are composed of control groups. Only one control group can be active per instruction fields during any clock cycle. The following instruction fields are defined:

```

instruction_field("IF_1",["CG_immediate_1"]); // 16 bits
instruction_field("IF_2",["CG_ram_1"]); // 2 bits
instruction_field("IF_3",["CG_alu_addsub_1"]); // 2 bits
25 instruction_field("IF_4",["CG_immediate_1","CG_alu_full_1"]); // 1+ 16 bits

```

Instruction sets are defined as a concatenation of instruction fields. The following instruction sets are defined:

```

30 instruction_set("IS_VLIW",["IF_1","IF_2","IF_3"]); // 2 + ( 2 + 2 + 16 ) bits
instruction_set("IS_RISC1",["IF_2","IF_4"]); // 2 + ( 1 + 16 ) bits
instruction_set("IS_RISC2",["IF_2","IF_3"]); // 2 + ( 2 ) bits

```

N = 4 bits

M = 5 banks

IS\_VLIW controls all datapath elements in parallel. IS\_RISC1 controls the ram and  
5 either the constant generator or the full functionality of the alu. IS\_RISC2 controls the  
the ram and a subset of the alu functionality.

The computer programs for controlling VLIW processors and including mixed  
instruction set instructions in accordance with the above techniques represent another  
10 aspect of this invention. These computer programs may be provided in the form of  
products such as data carriers (e.g. disks, ROMs, etc) or as downloadable files or in  
other ways.

With regard to compiler support instruction sets are stored in a data-structure  
15 that explicitly models allowed parallel operations in the set, or its dual, conflicting  
operations. Operations are conflicting when they use the same resources to operate on  
non-identical data sets. This data-structure can be used by a compiler for instruction  
selection and code generation. The instruction words created by the compiler are then  
selected from any of the defined instruction sets. In one example, this conflict model  
20 can be used in a Map step, a Compile step and a Code Generation step. Map selects  
instructions and assigns them to individual data-path elements. Compile schedules  
instructions and performs code optimization. At the end of Compile, code is emitted  
and final instruction set selection takes place. During this code generation step the  
jump addresses in the code are re-coded to allow for the variable length instruction  
25 words, which are stored distributed over multiple memories. In all three steps the cost  
of instruction sets (in terms of size) is taken into account.



**CLAIMS**

1. Apparatus for processing data, said apparatus comprising:  
an instruction decoder operable to decode one or more instruction fields within  
5 a program instruction to generate at least control signals; and  
a plurality of data path elements responsive to said controls signals  
independently to perform in parallel respective data processing operations specified  
by said program instruction, at least some program instructions decoded by said  
instruction decoder allowing control signals for respective data path elements to be  
10 independently specified; wherein  
said instruction decoder is operable such that a given data processing operation  
to be performed by a data path element can be specified using differently encoded  
instruction fields in different program instructions.
- 15 2. Apparatus as claimed in claim 1, wherein said differently encoded instruction  
fields representing said data processing operation have a different bit length within  
said different program instructions.
3. Apparatus as claimed in any one of claims 1 and 2, wherein said differently  
20 encoded instruction fields representing said data processing operation have a different  
bit position within said different program instructions.
4. Apparatus as claimed in any one of claims 1, 2 and 3, wherein said program  
instructions are divided into a plurality of instruction sets and said instruction decoder  
25 is responsive to at least an instruction set identifying field within a program  
instruction to control which bits within said program instruction correspond to  
instruction fields for which data path elements.
5. Apparatus as claimed in any one of the preceding claims, wherein said  
30 program instructions have a variable program instruction length and said instruction  
decoder is responsive to said instruction set identifier to determine program  
instruction length of a program instruction.

6. Apparatus as claimed in any one of the preceding claims, wherein said differently encoded instruction fields are operable to represent different sets of data processing operations that can be performed by said data path element.

5

7. Apparatus as claimed in claim 6, wherein at least one of said differently encoded instruction fields within a program instruction is operable to represent a subset of data processing operations that can be represented by a different one of said differently encoded instruction fields within a different program instruction.

10

8. Apparatus as claimed in claim 7, wherein said subset has a reduced operand range compared with said maximum set of data processing operations.

9. Apparatus as claimed in claim 8, wherein said operand range may be one of:

15

a register specifier;

a jump address specifier representing a jump in execution point within a program;

an alignment specifier; and

an immediate value.

20

10. Apparatus as claimed in any one of the preceding claims, wherein at least one instruction field within a program instruction itself specifies which of said plurality of data path elements is controlled by said at least one instruction field.

25 11. Apparatus as claimed in any one of the preceding claims, wherein in response to different program instructions variable numbers of said plurality of data path elements are active or inactive.

30 12. Apparatus as claimed in claim 11, wherein if a data processing operation is not specified or pending for a given data path element, then said given data path element is inactive and performs a Nop operation.

13. Apparatus as claimed in claim 5, comprising an instruction memory operable to store said program instructions and readable in fixed length memory accesses.

14. Apparatus as claimed in claim 13, wherein a fixed length memory access  
5 contains at least portions of a plurality of program instructions.

15. Apparatus as claimed in claim 14, comprising a rotator and a program counter incrementer operable to align program instructions within said memory accesses to a predetermined position for supply to said instruction decoder.

10

16. Apparatus as claimed in claim 5 and claim 15, wherein said rotator is responsive to said instruction set identifiers to determine what rotation to apply to align said program instructions.

15 17. A method of processing data, said method comprising the steps of:  
decoding one or more instruction fields within a program instruction to generate at least control signals; and

in response to said controls signals controlling a plurality of data path elements independently to perform in parallel respective data processing operations  
20 specified by said program instruction, at least some program instructions allowing control signals for respective data path elements to be independently specified; wherein

a given data processing operation to be performed by a data path element can be specified using differently encoded instruction fields in different program  
25 instructions.

18. A method as claimed in claim 17, wherein said differently encoded instruction fields representing said data processing operation have a different bit length within said different program instructions.

30

19. A method as claimed in any one of claims 17 and 18, wherein said differently encoded instruction fields representing said data processing operation have a different bit position within said different program instructions.

5 20. A method as claimed in any one of claims 17, 18 and 19, wherein said program instructions are divided into a plurality of instruction sets and said instruction decoder is responsive to at least an instruction set identifying field within a program instruction to control which bits within said program instruction correspond to instruction fields for which data path elements.

10

21. A method as claimed in any one of claims 17 to 20, wherein said program instructions have a variable program instruction length and said instruction decoder is responsive to said instruction set identifier to determine program instruction length of a program instruction.

15

22. A method as claimed in any one of claims 17 to 21, wherein said differently encoded instruction fields are operable to represent different sets of data processing operations that can be performed by said data path element.

20 23. A method as claimed in claim 22, wherein at least one of said differently encoded instruction fields within a program instruction is operable to represent a subset of data processing operations that can be represented by a different one of said differently encoded instruction fields within a different program instruction.

25 24. A method as claimed in claim 23, wherein said subset has a reduced operand range compared with said maximum set of data processing operations.

25. A method as claimed in claim 24, wherein said operand range may be one of:  
a register specifier;

30 a jump address specifier representing a jump in execution point within a program;  
an alignment specifier;and

an immediate value.

26. A method as claimed in any one of claims 17 to 25, wherein at least one instruction field within a program instruction itself specifies which of said plurality of data path elements is controlled by said at least one instruction field.

27. A method as claimed in any one of claims 17 to 26, wherein in response to different program instructions variable numbers of said plurality of data path elements are active or inactive.

10

28. A method as claimed in claim 27, wherein if a data processing operation is not specified or pending for a given data path element, then said given data path element is inactive and performs a Nop operation.

15 29. A method as claimed in claim 21, wherein an instruction memory is operable to store said program instructions and readable in fixed length memory accesses.

30. A method as claimed in claim 29, wherein a fixed length memory access contains at least portions of a plurality of program instructions.

20

31. A method as claimed in claim 30, wherein a rotator and a program counter incrementer are operable to align program instructions within said memory accesses to a predetermined position for supply to said instruction decoder.

25 32. A method as claimed in claim 21 and claim 31, wherein said rotator is responsive to said instruction set identifiers to determine what rotation to apply to align said program instructions.

30 33. A computer program product for controlling an apparatus for processing data having an instruction decoder operable to decode one or more instruction fields within a program instruction to generate at least control signals and a plurality of data path elements responsive to said controls signals independently to perform in parallel

respective data processing operations specified by said program instruction; said computer program product comprising:

at least one program instruction including an instruction field with a first encoding specifying a given data processing operation to be performed by a data path element; and

at least one different program instruction including a differently encoded instruction field with a second encoding different to said first encoding and specifying said given data processing operation to be performed by said data path element.

34. A computer program product as claimed in claim 33, wherein said differently encoded instruction fields representing said data processing operation have a different bit length within said different program instructions.

35. A computer program product as claimed in any one of claims 33 and 34, wherein said differently encoded instruction fields representing said data processing operation have a different bit position within said different program instructions.

36. A computer program product as claimed in any one of claims 33, 34 and 35, wherein said program instructions are divided into a plurality of instruction sets and said instruction decoder is responsive to at least an instruction set identifying field within a program instruction to control which bits within said program instruction correspond to instruction fields for which data path elements.

37. A computer program product as claimed in any one of claims 33 to 36, wherein said program instructions have a variable program instruction length and said instruction decoder is responsive to said instruction set identifier to determine program instruction length of a program instruction.

38. A computer program product as claimed in any one of claims 33 to 37, wherein said differently encoded instruction fields are operable to represent different sets of data processing operations that can be performed by said data path element.

39. A computer program product as claimed in claim 38, wherein at least one of said differently encoded instruction fields within a program instruction is operable to represent a subset of data processing operations that can be represented by a different one of said differently encoded instruction fields within a different program instruction.

40. A computer program product as claimed in claim 39, wherein said subset has a reduced operand range compared with said maximum set of data processing operations.

41. A computer program product as claimed in claim 40, wherein said operand range may be one of:

a register specifier;

a jump address specifier representing a jump in execution point within a program;

an alignment specifier; and

an immediate value.

42. A computer program product as claimed in any one of claims 33 to 41, wherein at least one instruction field within a program instruction itself specifies which of said plurality of data path elements is controlled by said at least one instruction field.

43. A computer program product as claimed in any one of claims 33 to 42, wherein in response to different program instructions variable numbers of said plurality of data path elements are active or inactive.

44. A computer program product as claimed in claim 43, wherein if a data processing operation is not specified or pending for a given data path element, then said given data path element is inactive and performs a null operation.

45. A computer program product as claimed in claim 37, wherein an instruction memory is operable to store said program instructions and readable in fixed length memory accesses.

5 46. A computer program product as claimed in claim 45, wherein a fixed length memory access contains at least portions of a plurality of program instructions.

47. A computer program product as claimed in claim 46, wherein a rotator and a program counter incrementer are operable to align program instructions within said  
10 memory accesses to a predetermined position for supply to said instruction decoder.

48. A computer program product as claimed in claim 37 and claim 47, wherein said rotator is responsive to said instruction set identifiers to determine what rotation to apply to align said program instructions.

15

49. A computer program compiler operable to generate a computer program as claimed in any one of claims 33 to 48.

50. Apparatus for processing data substantially as hereinbefore described with  
20 reference to the accompanying drawings.

51. A method of processing data substantially as hereinbefore described with reference to the accompanying drawings.

25 52. A computer program product substantially as hereinbefore described with reference to the accompanying drawings.

53. A computer program compiler substantially as hereinbefore described with reference to the accompanying drawings.





INVESTOR IN PEOPLE

Application No: GB0410986.4

Examiner: David Midgley

Claims searched: 1-53

Date of search: 13 September 2004

## Patents Act 1977: Search Report under Section 17

### Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
A	1,17,33	US5057837 A COLWELL

### Categories:

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

### Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC<sup>W</sup> :

G4A

Worldwide search of patent documents classified in the following areas of the IPC<sup>07</sup>

G06F

The following online and other databases have been used in the preparation of this search report

Online: WPI, EPODOC, PAJ/JAPIO, TDB, INSPEC, XPESP