

(19) World Intellectual Property
Organization
International Bureau




(43) International Publication Date
13 May 2004 (13.05.2004)

PCT

(10) International Publication Number
WO 2004/040445 A1

- (51) International Patent Classification⁷: **G06F 9/45**
 - (21) International Application Number:
PCT/RU2002/000469
 - (22) International Filing Date: 29 October 2002 (29.10.2002)
 - (25) Filing Language: English
 - (26) Publication Language: English
 - (71) Applicant (for all designated States except US): **MOTOROLA, INC.** [US/US]; 1303 East Algonquin Road, Schaumburg, IL 60196 (US).
 - (72) Inventors; and
 - (75) Inventors/Applicants (for US only): **OWENS, Howard, Dewey** [US/US]; 1808 Romeria Drive, Austin, TX 78757-3326 (US). **KIRILLIN, Vyacheslav Alexeevich** [RU/RU]; ul. Marshala Kazakova, 73-12, St.Petersburg, 198302 (RU). **KUTUZOV, Mikhail Andreevich** [RU/RU]; ul. Kompozitorov, 38-11/1, St.Petersburg, 194355 (RU). **PREOBRAZHENSKY, Dmitry Sergeevich** [RU/RU]; ul. Kolpinskaya, 15-17, St.Petersburg, 197110 (RU).
 - (74) Agents: **EGOROVA, Galina** et al.; Law Firm "Gorodissky & Partners" LTD, ul. Bolshaya Spasskaya, 25-3, Moscow, 129010 (RU).
 - (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.
 - (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— with international search report
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

(54) Title: METHOD AND APPARATUS FOR SELECTIVELY OPTIMIZING INTERPRETED LANGUAGE CODE



(57) Abstract: In one embodiment of the present invention an interpreted language, such as, for example, Java, is selectively optimized by partitioning the interpreted language code (98) into a plurality of blocks (80-83) based on the complexity of each of the interpreted language instructions. In one embodiment of the present invention, each of the plurality of blocks is identified as either a block to be compiled into native code (80-82) if the block is simple, or a block to be interpreted (83) if the block is complex. The compiled and interpreted blocks are appended to form in-line mixed code (99) that contains both native code (90-92) and interpreted language code (93). This mixed code is formed before run-time, so that no further compilation is required at run-time. A processing unit (102) may be used to execute the native code directly without the use of the Java VM (10), while also executing, in-line, the interpreted language code (93) which requires use of the Java VM (10) to interpret the Java bytecodes.

WO 2004/040445 A1

METHOD AND APPARATUS FOR SELECTIVELY OPTIMIZING INTERPRETED LANGUAGE CODE

Field of the Invention

The present invention relates generally to an
5 interpreted language, and more specifically to a method and
apparatus for selectively optimizing an interpreted
language.

Background of the Invention

Interpreted languages, such as Java, are high-level
10 languages compiled to an intermediate level that requires
an extra level of indirection to execute. For example, an
interpreted language, such as Java, is independent of the
hardware platform. It is generally more difficult for
software code written in an interpreted language to breach
15 the security of the host system which is executing the
intermediate level code.

Java is an object-oriented, multi-threaded programming
language that compiles to a compact intermediate form known
as bytecodes. Java is a common interpreted language which
20 is used to transfer applications over the internet.
Traditional Java technology generally cannot be efficiently
applied for embedded software development. Java bytecode
may be either directly executed by a Java bytecode
interpreter or accelerated by a Just-In-Time (JIT)
25 compiler. Both methods have their advantages and drawbacks.
Java bytecode interpreters require no or little memory for
execution, but the speed of interpretation is relatively
slow. Conventional JIT compilers are too big for embedded
applications and use a lot of memory, although JIT
30 compilers provide significant acceleration by compiling
Java bytecode into a native language at run-time. Some
compilers minimize the resources used by compiling only
performance-crucial fragments of applications. However,

these compiled fragments may contain relatively complex instructions (e.g. method invocation instructions), which results in huge generated code and additional overhead for compilation.

5 Brief Description of the Drawings

The present invention is illustrated by way of example and not limited by the accompanying figures, in which like references indicate similar elements, and in which:

10 FIG.1 illustrates one embodiment of a Java virtual machine 10 in accordance with one embodiment of the present invention;

 FIG.2 illustrates one embodiment of the Java method optimization 16 of FIG.1 in accordance with one embodiment of the present invention;

15 FIG.3 illustrates one embodiment of step 20 of FIG.2 in which blocks to be compiled or interpreted are identified in accordance with one embodiment of the present invention;

 FIG.4 illustrates one embodiment of step 21 of FIG.2 in which selected blocks identified in step 20 are now compiled in accordance with one embodiment of the present invention;

20 FIG.5 illustrates one embodiment of Java bytecode interpreter 17 and interpreter extension 18 of FIG.1 in accordance with one embodiment of the present invention;

 FIG.6 illustrates one embodiment of the code 98 (i.e. original bytecode for Java VM 10) and a translation which produced code 99 (i.e. mixed code structure for Java VM 10); and

30 FIG.7 illustrates a data processing system 100 in accordance with one embodiment of the present invention.

Skilled artisans appreciate that elements in the figures are illustrated for simplicity and clarity and have

not necessarily been drawn to scale. For example, the dimensions of some of the elements in the figures may be exaggerated relative to other elements to help improve the understanding of the embodiments of the present invention.

5

Detailed Description

In one embodiment of the present invention, an interpreted language (e.g. Java) is selectively optimized by partitioning the interpreted language code into a plurality of blocks based on the complexity of each of the interpreted language instructions. In one embodiment of the present invention, each of the plurality of blocks is identified as either a block to be compiled into native code if the block is simple, or a block to be interpreted (e.g. left as Java bytecodes) if the block is complex. In one embodiment of the present invention, a simple instruction is a Java bytecode that does not have any dependencies on Java VM 10 services (e.g. memory allocation, garbage collection, etc.). The compiled and non-compiled (i.e. interpreted) blocks are appended to form in-line mixed code (e.g. 99 in FIG.6) that contains both native code (e.g. 90-92) and interpreted language code (e.g. 93). It is this in-line mixed code that is executed at run time. Thus, no JIT compiler is required at run time. A processing unit (e.g. 102 in FIG.7) may be used to execute the native code directly without the use of a Java VM 10, while also executing, in-line, the interpreted language code which requires use of the Java VM 10 to perform the interpretation of the Java bytecodes. Consequently, for simple blocks, the extra level of indirection added by the Java VM 10 can be avoided, thus saving time and/or memory for systems (e.g. 100 in FIG.7) which are executing an interpreted language. This time/memory savings may be especially important for

portable or handheld devices which can download files from the internet, and thus can execute an interpreted language such as Java.

As used herein, the term "bus" is used to refer to a
5 plurality of signals or conductors which may be used to transfer one or more various types of information, such as data, addresses, control, or status.

FIG.1 illustrates one embodiment of a Java virtual machine 10 in accordance with one embodiment of the present
10 invention. In one embodiment, the present invention utilizes a Java virtual machine (VM) 10 which receives Java class files 12 from a source external to the Java VM 10. The Java VM 10 includes a class loader 14 which loads one or more Java classes from Java class files 12. The class
15 loader 14 provides Java class files to the portion of the Java VM 10 that is responsible for crucial method identification 15. The crucial method identification 15 process identifies performance-crucial functions (e.g. Java methods) of the loaded Java class files using a profiler or
20 externally-supplied information, e.g. special Java method attributes. Any appropriate process for performing the crucial method identification 15 process may be used (e.g. profiling). For some applications, 80% of the application's execution time is spent executing 20% of the application's
25 code. Thus acceleration of this 20% of the application's code may very significantly improve performance. The performance-crucial functions are passed from step 15 to step 16 where the Java methods are optimized. The method optimization step 16 is described in more detail in FIG.2.

30 The output of method optimization step 16 is mixed code 19 which may include both interpreted language instructions (e.g. bytecodes for Java) and native instructions. In one embodiment of the present invention,

the interpreter extension 18 extends the original Java
bytecode interpreter 17 by recognizing and handling a
special instruction (i.e. Java bytecode) called
"run_native" which transfers control from the interpreted
5 language instructions to the native instructions of the
mixed code. Control of the host processor (e.g. 102 of
FIG.7) is transferred from executing the Java virtual
machine 10 to executing a subsequent native instruction in
the mixed code. In one embodiment, the interpreter
10 extension 18 implements an efficient binary interface with
the compiled native code; for example, the interpreter
extension 18 may cache the most significant variables of
the Java bytecode interpreter 17 in registers. When the
interpreter extension 18 encounters a run_native
15 instruction, it transfers control to the following compiled
native code. Other Java bytecode instructions are
interpreted using the Java bytecode interpreter 17. The
interpreter extension 18 may be implemented in any manner.

In one embodiment of the present invention, to
20 minimize Java operand stack access in the native code,
stack values used by the native code are cached in a
special register file (e.g. in processing unit 102 or
memory 104 of FIG.7). The register file is a set of
registers which mimic top Java operand stack values. In one
25 embodiment of the present invention, in the beginning of
the compiled code, all Java stack values used by the
compiled code are transferred into the register file. At
the end of the compiled code, all new values are copied
back into the Java operand stack. The size of the compiled
30 code is limited so that all used stack values are kept in
the register file. Fortunately, due to the nature of Java
applications, most compiled fragments use five or less Java
stack elements. Moreover, most compiled fragments do not

transfer any data to their neighbors via the Java operand stack so the generated binary code is usually very compact.

FIG.2 illustrates one embodiment of the Java method optimization of FIG.1 in accordance with one embodiment of the present invention. In step 20, the blocks to be compiled are identified and the blocks to be interpreted are identified. Then in step 21, the blocks selected to be compiled are actually compiled. Step 22 resolves any jump references in the mixed code, and step 23 modifies the exception table as necessary.

The present invention thus compiles only the most profitable blocks of the Java method's bytecodes which may be significantly accelerated without much memory overhead; other blocks are left untouched. The set of compiled instructions depends on the architecture of the Java bytecode interpreter 17 and the target processor (e.g. processing unit 102 of FIG.7). In one embodiment of the present invention, the blocks are selected so that each block has just one entry point and one exit point. In some embodiments, native code for each compiled Java instruction will not exceed 10-15 instructions of the target processor and will not include subroutine calls. Alternate embodiments of the present invention may set other limits on determining which instructions will be compiled. Note that compilation of complex instructions does not necessarily improve performance, but takes up additional resources and complicates optimization. The compiled blocks of Java bytecode are replaced by the generated binary code prefixed with a special bytecode instruction `run_native`.

"Run_native" is a predetermined interpreted language instruction which indicates to the interpreter extension 18 (see FIG.1) that the following code is native code. The resultant mixed code 19 from method optimization 16

consists of blocks of native code and Java bytecode instructions which cannot be well accelerated. The mixed code structure of a Java method having mixed code is illustrated in Figure 6. Alternate embodiments of the present invention may use other approaches to indicate to the interpreter extension that the following or subsequent in-line code is native code.

FIG.3 illustrates one embodiment of step 20 of FIG.2 in which blocks to be compiled or interpreted are identified in accordance with one embodiment of the present invention. In step 30, jump targets are identified. The flow then continues to step 31 where the variable "i" is set to zero. The variable "i" indicates which bytecode is currently being processed. In step 32, the current bytecode "bc" is set equal to bytecode(i). Note that alternate embodiments of the present invention may perform step 32 as the first step in the "NO" path after decision diamond 38 with step 31 linked directly with step 32 and step 34 being the input for decision diamond 38. At decision diamond 38 the question is asked "is the current bytecode "bc" the last instruction?". If "bc" is the last instruction, then the end has been reached and the process continues with step 21 in FIG.2. If "bc" is not the last instruction, then the process continues to decision diamond 39 where the question is asked "is bc a jump target?". If bc is a jump target, the process continues to step 37 where the current block to be compiled (if any) is finished, and the process continues to decision diamond 41. If bc is not a jump target, the process continues to decision diamond 41 where the question is asked " is bc a simple instruction?"

In one embodiment of the present invention, a simple instruction is a Java bytecode that does not have any dependencies on Java VM 10 services (e.g. memory

allocation, garbage collection, etc.). Alternate embodiments of the present invention may use any desired criteria to determine which interpreted language instructions are "simple". If the current bytecode "bc" is not a simple instruction, then the process continues to step 35 where the current block to be compiled, if there is any, is finished. From step 35, the process continues to step 34 where the process determines the length of the current bytecode "bc" so that the flow can move to the beginning of the next bytecode. From step 34, the process continues to step 32 where the next bytecode becomes the current bytecode. If the current bytecode is a simple instruction, then the process continues to decision diamond 40 where the question is asked " is there a current block to be compiled?". If there is a current block to be compiled, the process continues to step 33 where the current bytecode is added to the current block to be compiled. If there is not a current block to be compiled, the process continues to step 36 where a new block to be compiled is created. From step 36, the process continues to step 33 where the next bytecode becomes the current bytecode (e.g. by adding "bc" to the current block to be compiled). The process then continues to step 34, then step 32, then decision diamond 38 as described above.

FIG.4 illustrates one embodiment of step 21 of FIG.2 in which selected blocks identified in step 20 are now compiled in accordance with one embodiment of the present invention. The process starts at decision diamond 50 where the question is asked "is the current block a block to be compiled?".

If the current block is a block to be compiled, i.e. is a block to be compiled as native code, the process continues at step 55 where the variable "i" is set equal to

the block offset. The process continues to step 56 where "bc" is set equal to the current bytecode "bytecode(i)". The process continues to step 57 where the current bytecode "bc" is compiled. The compilation step 57 results in
5 compiled code in the native language of processing unit 102 (see FIG.7). The process continues to step 58 where the compiled code in the native language is appended to the mixed code (see right-hand column in FIG.6). The process continues to step 59 where the process determines the
10 length of the current bytecode so that the flow can move to the beginning of the next bytecode. The process continues to decision diamond 51 where the question is asked "is bc the last instruction in the block?". If the current bytecode is the last one in the block, the process
15 continues to step 22 in FIG.2. If the current bytecode is not the last one in the block, the process continues to step 56 where the next bytecode in the block is made the current bytecode, and the steps 57-59 are repeated.

If the current block is not a block to be compiled, 20 i.e. is a block to be left as Java bytecodes, then the flow continues from decision diamond 50 to step 52 where a native header (e.g. native header 95 in FIG.6) is appended to the mixed code. The native header 95 can be used to return control from the native code back to the Java
25 bytecode interpreter 17 (see FIG.1). The mixed code may include both Java bytecodes and native instructions. From step 52, the process continues to step 53 where the interpreted bytecode is appended to the mixed code. In one embodiment of the present invention, no compilation of the
30 interpreted bytecodes is performed. The interpreted bytecodes remain unchanged and are merely appended as they are to the in-line mixed code 99 of FIG.6. From step 53, the process continues to step 54 where a special bytecode

called "run_native" is appended to the mixed code (see run_native 94 appended to mixed code 99 of FIG.6) in preparation of the next block to be compiled.

The special bytecode "run_native" is used to transfer
5 control from the interpreted language instructions to the native instruction of the mixed code. In one embodiment, control of the host processor (e.g. 102 of FIG.7) is transferred from executing the Java virtual machine 10 to executing a subsequent native instruction in the mixed
10 code.

In alternate embodiments of the present invention, steps 56-59 and decision diamond 51, which process one bytecode at a time, may be replaced by a parallel process that considers a plurality of bytecodes at a time in order
15 to perform further optimization.

FIG.5 illustrates one embodiment of Java bytecode interpreter 17 and interpreter extension 18 of FIG.1 in accordance with one embodiment of the present invention. At startup, Java bytecode interpreter 17 sets the current
20 bytecode "bc" equal to the next bytecode in step 70. The process continues to decision diamond 75 in interpreter extension 18 where the question is asked "is bc a special bytecode "run_native"?. If the current bytecode is not the special bytecode "run_native", the process continues to
25 step 74 where the Java bytecode interpreter 17 interprets the current bytecode. The process then continues back to step 70 where the next bytecode is selected. Returning to decision diamond 75, if the current bytecode is the special bytecode "run_native", the process continues
30 to step 73 where the next code to be executed is compiled code in native language. The process continues to step 72 where the compiled code in native language is executed. The process continues to step 71 where a return from the

compiled code is performed. In one embodiment of the present invention, the return from the compiled code is implemented by way of a native header 95 (see FIG.6). From step 71, the process continues to step 70 where the next
5 bytecode is selected. If the next bytecode is the last bytecode, and thus is of the type "return" at the highest level, then the processing unit 102 of FIG.7 stops executing the Java VM 10.

Note that the software used in the present invention
10 is not limited to the embodiments described in the flow diagrams. For example, the ordering of the steps and decision points described in the flow diagrams may be varied for different embodiments of the present invention. In addition, alternate embodiments of the present invention
15 may use different steps and/or decision diamonds than those illustrated in the flow diagrams.

FIG.6 illustrates one embodiment of the code 98 (i.e. original bytecode for Java VM 10) and a translation which produced code 99 (i.e. mixed code structure for Java VM
20 10). In one embodiment of the present invention, code 99 includes compiled code 90, 91, and 92 in-line with appended code 93. In one embodiment of the present invention, appended code 93 is Java bytecode. Run_native is a special
25 bytecode that is used as an instruction for the interpreter extension 18 (see FIG.1) to indicate that the following in-line code is native code to be executed directly by processing unit 102 (see FIG.7) without use of the Java VM 10. Native header 95 is used by processor 102 to return control back to the Java bytecode interpreter 17 within the
30 Java VM 10. Note that this mixed code 99 is formed before run-time, so that, unlike JIT compilers, no further compilation is required at run-time.

In one embodiment, code 98 is the original bytecode

for the Java VM 10. The specific instruction used in FIG.6 are for illustrative purposes only. Other instructions could have been used. Note that the interpreted language instructions (e.g. block 93) are still interpreted by the
5 Java VM 10 running on host processor 102, unlike the compiled code 90-92 which is native code that is executed directly by processor 102 without use of the Java VM 10.

Referring to code 98, Java bytecode "ILOAD_0" located at offset 0 and Java bytecode "IFLE_8" located at offset 1
10 together form a block 80 that is determined to be "simple" (see decision diamond 41 in FIG.3) and thus is to be compiled. Java bytecode "ILOAD_0" located at offset 4 and Java bytecode "GOTO 10" located at offset 5 together form a block 81 that is determined to be "simple" (see decision
15 diamond 41 in FIG.3) and thus is to be compiled. Java bytecode "ILOAD_0" located at offset 8 and Java bytecode "INEG" located at offset 9 together form a block 82 that is determined to be "simple" (see decision diamond 41 in FIG.3) and thus is to be compiled. Block 80 is compiled to
20 create compiled code 90; block 81 is compiled to create compiled code 91; and block 82 is compiled to create compiled code 92. A special bytecode instruction run_native 94 is placed in the in-line code just before the beginning of the blocks of compiled code 80-82. The native header 95
25 is placed in the in-line code at the end of the native code and just before the beginning of the interpreted code 93. In one embodiment of the present invention, the interpreted code 93 is the same as the block of code 83 that is to be interpreted. In one embodiment, the translation process
30 from code 98 to code 99 merely copies the original Java bytecodes from block 83 to block 93. In one embodiment, mixed code 99 now includes both compiled code in the native language and non-compiled code that is still Java

bytecodes. Code 98 and mixed code 99 may be stored in memory 104 (see FIG.7) or in any other portion of data processing system 100.

FIG.7 illustrates a data processing system 100 in accordance with one embodiment of the present invention. In one embodiment, data processing system 100 is a portable, handheld device. In one embodiment of the present invention, data processing system is capable of receiving information from the internet via information port 106. Although data processing system 100 has been shown to have the architecture illustrated in FIG.7, any architecture may be used for data processing system 100.

In one embodiment, data processing system 100 has a processing unit 102, a memory 104, an information port 106, other circuitry 108, and user interface 110 which are bi-directionally coupled to bus 116. In one embodiment of the present invention, memory 104 includes a Java virtual machine 10. In alternate embodiments of the present invention, Java VM 10 may be stored anywhere. Alternate embodiments of the present invention may use other circuitry 108 to implement any desired function. Alternate embodiments of data processing system 100 may not include information port 106, may not include other circuitry 108, and/or may not include user interface 110. User interface 110 may include anything which allows a user to communicate with data processing system 100, such as, for example, a keypad, a mouse, a display, a touch screen, or audio I/O.

In the foregoing specification, the invention has been described with reference to specific embodiments. However, one of ordinary skill in the art appreciates that various modifications and changes can be made without departing from the scope of the present invention as set forth in the claims below. For example, although various embodiments of

the present invention have been described in the context of Java, the present invention is applicable to any interpreted language, not just Java. Similarly, any architecture for data processing system 100 (see FIG.7) may
5 be used. Similarly, any software may be used to implement the claimed invention. Accordingly, the specification and figures are to be regarded in an illustrative rather than a restrictive sense, and all such modifications are intended to be included within the scope of present invention.

10 Benefits, other advantages, and solutions to problems have been described above with regard to specific embodiments. However, the benefits, advantages, solutions to problems, and any element(s) that may cause any benefit, advantage, or solution to occur or become more pronounced
15 are not to be construed as a critical, required, or essential feature or element of any or all the claims. As used herein, the terms "comprises," "comprising," or any other variation thereof, are intended to cover a non-exclusive inclusion, such that a process, method, article,
20 or apparatus that comprises a list of elements does not include only those elements but may include other elements not expressly listed or inherent to such process, method, article, or apparatus.

CLAIMS

1. In a virtual machine executing on a host processor, a method for selectively optimizing interpreted language code, comprising:
5 receiving interpreted language code comprising interpreted language instructions; and partitioning the interpreted language code into a plurality of blocks based on a complexity of each of the interpreted language instructions, each of
10 the plurality of blocks identified as one of a block to be compiled or a block to be interpreted.
2. The method of claim 1, wherein a first block of the plurality of blocks is identified as a block to be
15 compiled and a second block of the plurality of blocks is identified as a block to be interpreted, said method further comprising creating mixed code, said mixed code comprising native instructions and interpreted language instructions, wherein creating said mixed code
20 comprises:
compiling the first block;
appending the compiled first block to the mixed code, wherein the compiled first block includes a plurality of native instructions;
25 appending the second block to the mixed code, wherein the second block comprises a plurality of interpreted language instructions.
3. The method of claim 2, wherein the native instructions run on the host processor and the interpreted language
30 instructions are interpreted by the virtual machine.
4. The method of claim 2, wherein, in the mixed code, the interpreted language instructions are in-line with the native instructions.

5. The method of claim 2, wherein creating the mixed code further comprises:
- 5 appending a predetermined interpreted language instruction to the mixed code, wherein the predetermined interpreted language instruction transfers control from the interpreted language instructions to the native instructions of the mixed code.
6. The method of claim 5, wherein creating the mixed code further comprises:
- 10 appending a native header to the mixed code.
7. The method of claim 1, wherein for each of the interpreted language instructions in the interpreted language code, partitioning further comprises:
- 15 determining if the interpreted language instruction is a simple instruction;
- if the interpreted language instruction is a simple instruction, including the interpreted language instruction into one of the plurality of blocks identified as a block to be compiled; and
- 20 if the interpreted language instruction is not a simple instruction, including the interpreted language instruction into one of the plurality of blocks identified as a block to be interpreted.
- 25 8. The method of claim 7, wherein each block identified as a block to be compiled comprises no complex instructions.
9. The method of claim 7, wherein for each of the interpreted language instructions in the interpreted language code, partitioning further comprises:
- 30 determining if the interpreted language instruction is a jump target; and
- if the interpreted language instruction is a jump

target, finishing a current block within the plurality of blocks.

10. A data processing system for creating mixed code, said mixed code comprising native instructions and interpreted language instructions, said data processing system comprising:
- 5
- a processing unit for executing native instructions;
 - a memory coupled to the processing unit and having a virtual machine, wherein said virtual machine is executed by the processing unit and comprises:

10

 - a first set of instructions for receiving interpreted language code comprising interpreted language instructions;
 - a second set of instructions for partitioning the interpreted language code into a plurality of blocks based on a complexity of each of the interpreted language instructions, each of the plurality of blocks identified as one of a block to be compiled or a block to be interpreted;

15

 - a third set of instructions for compiling each of the plurality of blocks identified as a block to be compiled;
 - a fourth set of instructions for appending the compiled blocks to the mixed code, wherein the compiled blocks each comprises a plurality of native instructions; and

20

 - a fifth set of instructions for appending each of the plurality of blocks identified as a block to be interpreted to the mixed code, wherein each of the plurality of blocks identified as a block to be interpreted comprises a plurality of interpreted language
- 25
- 30

instructions.

11. The data processing system of claim 10, wherein the virtual machine further comprises:

5 a sixth set of instructions for receiving an interpreted language instruction from the mixed code;

10 a seventh set of instructions for determining if the received interpreted language instruction from the mixed code is a predetermined interpreted language instruction;

15 an eighth set of instructions for transferring control to a native instruction in the mixed code if the received interpreted language instruction from the mixed code is the predetermined interpreted language instruction; and

20 a ninth set of instructions for interpreting the received interpreted language instruction from the mixed code if the received interpreted language instruction from the mixed code is not the predetermined interpreted language instruction.

12. The data processing system of claim 10, wherein the second set of instructions further comprises:

25 a sixth set of instructions for determining if an interpreted language instruction is a simple instruction,

30 a seventh set of instructions for including the interpreted language instruction into one of the plurality of blocks identified as a block to be compiled if the interpreted language instruction is a simple instruction; and

an eighth set of instructions for including the interpreted language instruction into one of the plurality of blocks identified as a block to be

interpreted if the interpreted language instruction is not a simple instruction.

13. A hand held device comprising the data processing system of claim 10.

5 14. A data processing system for creating mixed code, said mixed code comprising native instructions and interpreted language instructions, said data processing system comprising:

10 means for receiving interpreted language code comprising interpreted language instructions;

means for partitioning the interpreted language code into a plurality of blocks based on a complexity of each of the interpreted language instructions,

15 each of the plurality of blocks identified as one of a block to be compiled or a block to be interpreted;

means for compiling each of the plurality of blocks identified as a block to be compiled;

20 means for appending the compiled blocks to the mixed code, wherein the compiled blocks each comprises a plurality of native instructions; and

25 means for appending each of the plurality of blocks identified as a block to be interpreted to the mixed code, wherein each of the plurality of blocks identified as a block to be interpreted comprises a plurality of interpreted language instructions.

15. The data processing system of claim 14, wherein the data processing system further comprises:

30 means for receiving an interpreted language instruction from the mixed code;

means for determining if the received interpreted language instruction from the mixed code is a

- predetermined interpreted language instruction;
means for transferring control to a native instruction
in the mixed code if the received interpreted
language instruction from the mixed code is the
5 predetermined interpreted language instruction;
and
means for interpreting the received interpreted
language instruction from the mixed code if the
received interpreted language instruction from the
10 mixed code is not the predetermined interpreted
language instruction.
16. The data processing system of claim 14, further
comprising:
means for determining if an interpreted language
15 instruction is a simple instruction,
means for including the interpreted language
instruction into one of the plurality of blocks
identified as a block to be compiled if the
interpreted language instruction is a simple
20 instruction; and
means for including the interpreted language
instruction into one of the plurality of blocks
identified as a block to be interpreted if the
interpreted language instruction is not a simple
25 instruction.
17. A hand held device comprising the data processing
system of claim 14.
18. A virtual machine stored on a computer readable medium,
said virtual machine capable of being executed by a
30 host processor, said virtual machine comprising:
a first set of instructions for receiving interpreted
language code comprising interpreted language
instructions;

- 5 a second set of instructions for partitioning the interpreted language code into a plurality of blocks based on a complexity of each of the interpreted language instructions, each of the plurality of blocks identified as one of a block to be compiled or a block to be interpreted;
- a third set of instructions for compiling each of the plurality of blocks identified as a block to be compiled;
- 10 a fourth set of instructions for appending the compiled blocks to the mixed code, wherein the compiled blocks each comprises a plurality of native instructions capable of being executed by the host processor; and
- 15 a fifth set of instructions for appending each of the plurality of blocks identified as a block to be interpreted to the mixed code, wherein each of the plurality of blocks identified as a block to be interpreted comprises a plurality of interpreted language instructions.
- 20 19. The virtual machine of claim 18, wherein the virtual machine further comprises:
- a sixth set of instructions for receiving an interpreted language instruction from the mixed code;
- 25 a seventh set of instructions for determining if the received interpreted language instruction from the mixed code is a predetermined interpreted language instruction;
- 30 an eighth set of instructions for transferring control to a native instruction in the mixed code if the received interpreted language instruction from the mixed code is the predetermined interpreted

- language instruction; and
- a ninth set of instructions for interpreting the received interpreted language instruction from the mixed code if the received interpreted language instruction from the mixed code is not the predetermined interpreted language instruction.
- 5
20. The virtual machine of claim 18, wherein the second set of instructions further comprises:
- a sixth set of instructions for determining if an interpreted language instruction is a simple instruction,
- 10
- a seventh set of instructions for including the interpreted language instruction into one of the plurality of blocks identified as a block to be compiled if the interpreted language instruction is a simple instruction; and
- 15
- an eighth set of instruction for including the interpreted language instruction into one of the plurality of blocks identified as a block to be interpreted if the interpreted language instruction is not a simple instruction.
- 20
21. In a virtual machine executing on a host processor, a method for executing mixed code, said mixed code comprising native instructions in-line with interpreted language instructions, the native instructions executed by the host processor and the interpreted language instructions interpreted by the virtual machine, said method comprising:
- 25
- receiving an interpreted language instruction from the mixed code;
- 30
- determining if the received interpreted language instruction from the mixed code is a predetermined interpreted language instruction which indicates

that the subsequent instruction is a native instruction;

transferring control of the host processor from executing the virtual machine to execute the subsequent instruction in the mixed code if the
5 received interpreted language instruction from the mixed code is the predetermined interpreted language instruction; and

interpreting the received interpreted language instruction from the mixed code if the received
10 interpreted language instruction from the mixed code is not the predetermined interpreted language instruction.

22. The method of claim 21, further comprising creating
15 said mixed code prior to receiving the interpreted language instruction from the mixed code, wherein creating said mixed code comprises:

receiving interpreted language code comprising interpreted language instructions; and
20 partitioning the interpreted language code into a plurality of blocks based on a complexity of each of the interpreted language instructions, each of the plurality of blocks identified as one of a block to be compiled or a block to be interpreted.




FIG.1




FIG.2




FIG.3




FIG.4

4/5

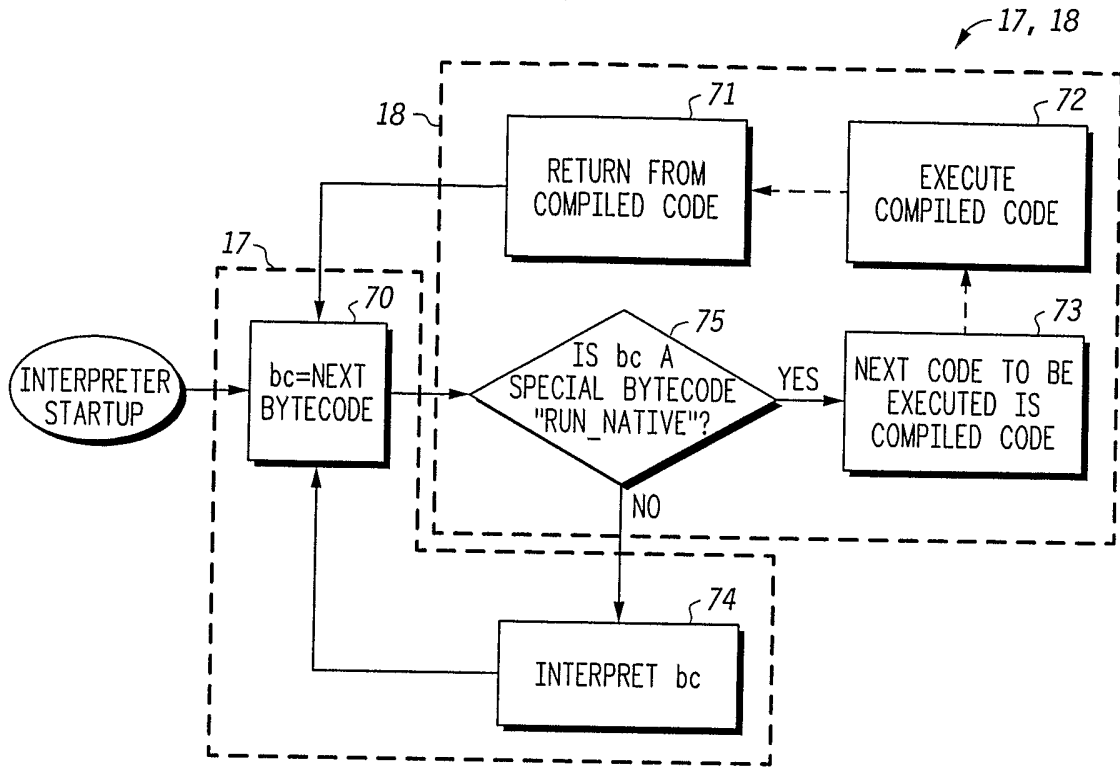


FIG.5

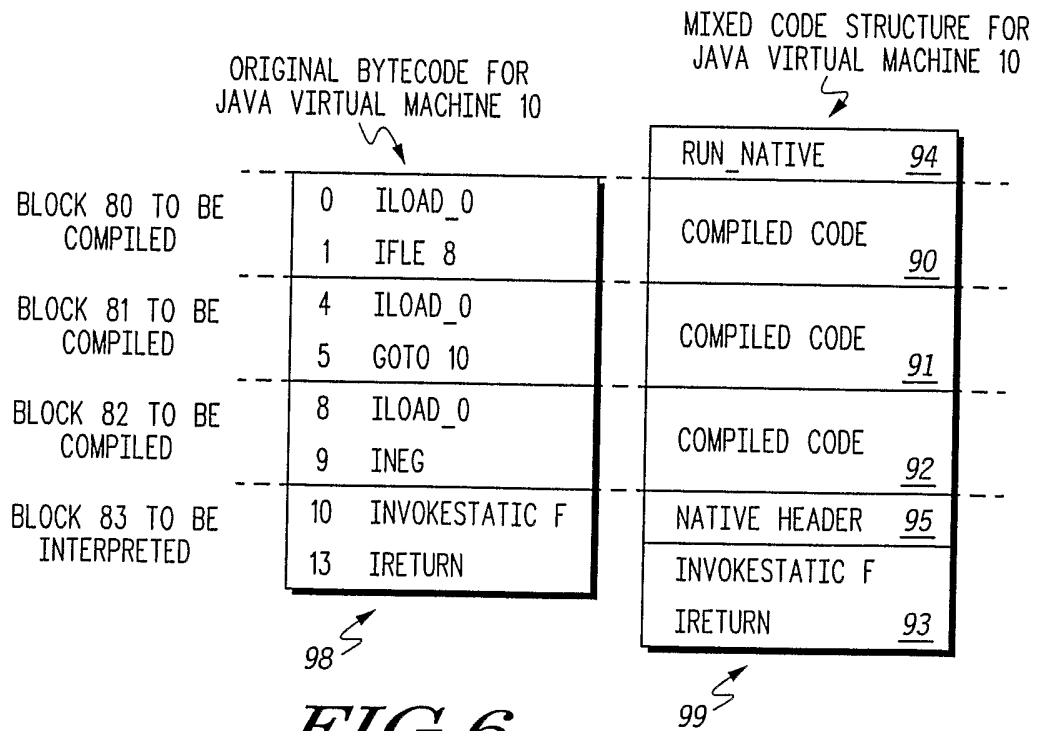


FIG.6

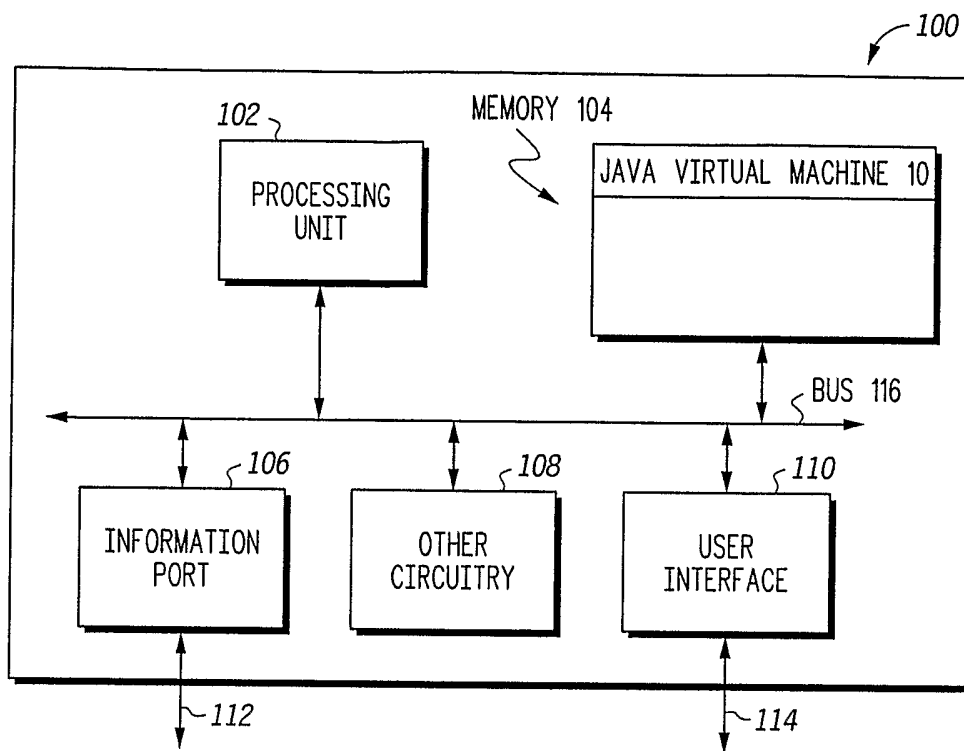


FIG. 7

INTERNATIONAL SEARCH REPORT

International application No.
PCT/RU 02/00469

A. CLASSIFICATION OF SUBJECT MATTER		G06F 9/45
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED		
Minimum documentation searched (classification system followed by classification symbols) G06F 9/00, 9/06, 9/30, 9/44, 9/45, 15/00, 15/02, 15/16, 13/00, H04L 9/00, 9/06, 9/28		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched:		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No
Y	US 6412109 B1 (SUN MICROSYSTEMS, INC.) Jun. 25, 2002, column 3, line 66-column 4, line 36	1
Y	US 6408433 B1 (SUN MICROSYSTEMS, INC.) Jun. 18, 2002, column 11, line 26-57, fig. 8	1
A	US 6412107 B1 (TEXAS INSTRUMENTS INCORPORATED) Jun. 25, 2002	1-22
A	RU 2147378 C1 (KOMMKVEST TEKNOLODZHIZ, INC.) 10.04.2000	1-22
<input type="checkbox"/> Further documents are listed in the continuation of Box C.		<input type="checkbox"/> See patent family annex
* Special categories of cited documents:		
"A"	document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E"	earlier document but published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L"	document with may throw doubts on priori claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O"	document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P"	document published prior to the international filing date but later than the priority date claimed	
Date of the actual completion of the international search report	28 May 2003 (28.05.2003)	Date of mailing of the international search report 05 June 2003 (05.06.2003)
Name and mailing address of the ISA/RU FIPS Russia, 123995, Moskva, G-59, GSP-5 Berezhkovskaya nab., 30-1	Authorized officer O. Krysanova	
Facsimile No.	Telephone No. (095)240-25-91	