

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
21 June 2007 (21.06.2007)

PCT

(10) International Publication Number
WO 2007/070073 A2

(51) International Patent Classification:
G06F 9/45 (2006.01)

Marcus [SE/SE]; Rorstrandsgatan 16, S-11340 Stockholm (SE).

(21) International Application Number:
PCT/US2006/012639

(74) Agents: MEYER, Sheldon, R. et al.; Fliesler Meyer LLP, 650 California Street, Fourteenth Floor, San Francisco, CA 94108 (US).

(22) International Filing Date: 5 April 2006 (05.04.2006)

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/748,763 9 December 2005 (09.12.2005) US
11/397,212 4 April 2006 (04.04.2006) US

(71) Applicant (for all designated States except US): BEA SYSTEMS, INC. [US/US]; 2315 North First Street, San Jose, California 95131 (US).

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),

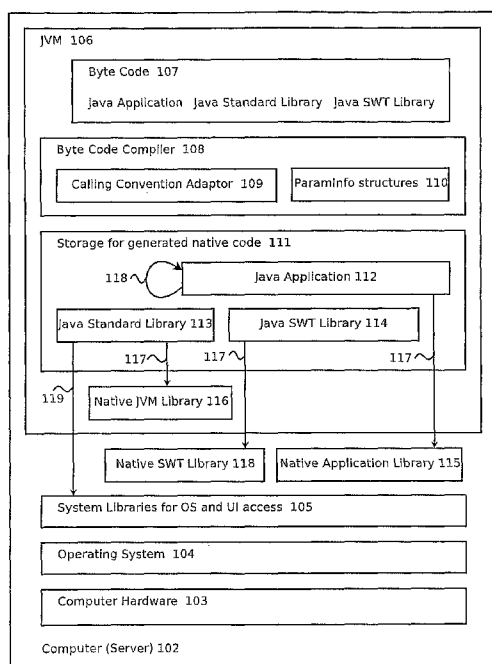
(72) Inventor; and

(75) Inventor/Applicant (for US only): LAGERGREN,

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR EFFICIENTLY GENERATING NATIVE CODE CALLS FROM BYTE CODE IN VIRTUAL MACHINES

(57) Abstract: A system and method for generation of efficient calls from compiled byte code within a virtual machine (VM) or Java Virtual Machine (JVM) environment. Current techniques used to interface dynamically compiled byte code with native libraries using stubs are unsatisfactory due to non-optimal performance and the large amount of manual work needed to maintain the stub generation code. An embodiment of the present invention addresses this by providing a system and a method for automatic stub generation based on a small set of parameters stored in a ParamInfo structure. In accordance with an embodiment, the system comprises a computer with native libraries; a virtual machine for executing a software application; a memory space for the compiled byte code; and a compiler with a calling convention adaptor that integrates native calls into the high level analysis step of the compilation of byte codes, so that the overhead for the native call is reduced.



WO 2007/070073 A2



European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

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.

Published:

- *without international search report and to be republished upon receipt of that report*

SYSTEM AND METHOD FOR EFFICIENTLY GENERATING NATIVE CODE CALLS FROM BYTE CODE IN VIRTUAL MACHINES

5 **Claim of Priority:**

U.S. Provisional Patent Application No. 60/748,763 entitled "SYSTEM AND METHOD FOR EFFICIENTLY GENERATING NATIVE CODE CALLS FROM BYTE CODE IN VIRTUAL MACHINES", by Marcus Lagergren *et al.*, filed December 9, 2005 (Attorney Docket No. BEAS-02013US0); and

10 U.S. Patent Application No. 11/____,____ entitled "SYSTEM AND METHOD FOR EFFICIENTLY GENERATING NATIVE CODE CALLS FROM BYTE CODE IN VIRTUAL MACHINES", by Marcus Lagergren, filed April 4, 2006 (Attorney Docket No.: BEAS-02013US1).

15 **Field of the Invention:**

The invention is generally related to virtual machine environments and other run-time environments, and particularly to a system and method for improving generation of native code calls from byte code.

20 **Background:**

Virtual machines are abstract computers that allow for portability of software applications, typically between different underlying computer architectures. A virtual machine (VM) is generally a complex software product that is implemented upon a particular computer hardware platform and/or operating system. The VM
25 then provides a uniform layer of abstraction between the hardware platform and any compiled software applications that will run thereon. Virtual machines are essential for the portability of certain technologies, including Java. The Java Virtual Machine (JVM) allows compiled Java programs to be run on the JVM, independently of whatever hardware or operating system may be underneath.
30 The JVM is described in further detail in the book "The Java™ Virtual Machine Specification (2nd Edition)" by Tim Lindholm, published by Sun Microsystems, and

incorporated herein by reference. Examples of commercially available JVMs include the Sun Java Virtual Machine from Sun Microsystems, Inc., and the JRockit Virtual Machine from BEA Systems, Inc.

5 A real CPU understands and executes instructions native to that CPU (commonly called native code) in comparison a virtual machine understands and executes virtual machine instructions (commonly called byte code). A virtual machine almost always run on a real CPU executing native code. The core of a virtual machine is normally implemented in a programming language such as C, that is always compiled to native code using an OS/CPU compatible compiler.

10 A virtual machine can implement different strategies of how to execute the byte codes. If the virtual machine analyzes each byte code separately and does this every time the same byte code is executed, then the the virtual machine is said to be an "interpreter". If the virtual machine translates the byte code into native code once and then the native code is used every time the same byte code
15 is executed, then the virtual machine is said to be a "just in time compiler" (commonly called a JIT).

Some virtual machines contain both an interpreter and a JIT. In the case of Java Virtual Machines, the Sun Java Virtual Machine will initially use the interpreter when executing Java byte code. When the Sun JVM detects byte code
20 that is executed often (commonly called a hot spot in the program), then it will compile that part of the byte code into native code. By contrast, the JRockit Virtual Machine from BEA will never interpret the Java byte code. It will always compile it to native code before executing it. If JRockit detects a hot spot in the program it will recompile that part of the byte code again, but with more code optimizations.

25 A Java Virtual Machine always needs to call native code to access operating system resources. Since the core of the Java Virtual Machine is written in a language such as C that can be compiled and linked to the operating system libraries, accessing operating system resources is simply a native function call following the platform calling conventions.

30 A JIT tries to optimize function calls between Java functions to use the most efficient way of calling on the particular CPU. One way to do this is to use

registers as much as possible for arguments to functions. Due to several reasons this is usually not the same as the platform calling convention. For example, a JIT for a language with garbage collecting (like Java) needs to take care which registers contain pointers to live objects, the Java-to-Java calling convention can therefore declare that certain registers always contain object pointers and not temporary results from arithmetic calculations. Such care is not needed for the native calling convention. On a register-starved architecture like the Intel x86 processors, the JIT can also use fewer callee save registers than the platform calling convention and instead use the remaining registers for function arguments.

However, often byte code needs to make use of native libraries. These libraries can be used for a variety of purposes, including: low level graphics (like the implementation of the Standard Widget Library (SWT) from the Eclipse project); database access (native drivers are sometimes required to speed up access to databases); or large amounts of legacy code that cannot be ported.

In the Java language this is solved using the standardized Java Native Interface (JNI). The JNI specifies that the native code should be called using the platform calling convention, and also specifies how the Java arguments are translated to a format that native code can use. JNI is described in further detail in the book "Java™ Native Interface: Programmer's Guide and Specification" by Sheng Liang, which is incorporated herein by reference.

Traditionally, a call from compiled Java code to native code is redirected through a short piece of native code (a stub), and the stubs are generated by a stub generator. The stub performs the translation of the arguments, sets up the call to conform to the platform calling convention and finally calls the native code. Depending on the JIT strategy, some arguments, such as numbers need not be translated. In those cases the overhead introduced by the extra function calls are more noticeable than when several necessary argument translations are part of the call. If the native function that is called is very short, then the overhead of the native call setup can be significant.

Platform/Environment	A Java call	A native call
x86_64/Windows	Ptr args: rsi,rdi Inte args: rax,rdx Stack used for more args. Calle save: rbx,rbp	Args: rcx,rdx,r8,r9 Stack used for more args.
x86_64/Linux	Same as above	Args: rdi,rsi,rcx,rdx,r8,r9 Stack used for more args.
x86/Windows & Linux	Ptr args: esi,edi Inte args: eax,edx Stack used for more args. Calle save: ebx,ebp	Stack used for all args.
ia64/Windows & Linux	Same as right.	Variable sized stack frame on register stack. Args: r32,r33,...

Table 1

It is also a significant amount of work to write the stub generators given the number of common operating systems, for example AIX, Linux, Solaris and Windows NT/2k/XP, and their respective calling conventions. Table 1 shows that Windows NT/2k/XP and Linux use the same calling convention on Intel x86 and ia64 processors. However Windows and Linux use different calling conventions on Intel EM64T compatible processors (x86_64). Other calling conventions include AIX on IBM PowerPC processors, and Solaris on SPARC processors.

```

public class HelloWorld
{
    native void printHelloWorld(int a);
    static { System.loadLibrary("nativelib"); }
    public static void main(String[] args) {
        HelloWorld h = new HelloWorld();
        h.helloWorld(0x1111);
        h.helloWorld(0x2222);
        h.helloWorld(0x3333);
        h.helloWorld(0x4444);
        h.helloWorld(0x5555);
        h.helloWorld(0x6666);
    }
}

```

Listing 1

Listing 1 demonstrates an example of a Java-program that calls a native library. A class-file, such as that shown in Listing 2, is generated when the

program is compiled with javac.

```

0:   new      #2; //class HelloWorld
3:   dup
5:   4:   invokespecial  #3; //Method "<init>":()V
7:   astore_1
8:   aload_1
9:   sipush  4369
10:  12:  invokevirtual#4; //Method printHelloWorld:(I)V
15:  15:  aload_1
16:  16:  sipush  8738
19:  19:  invokevirtual#4; //Method printHelloWorld:(I)V
22:  22:  aload_1
23:  23:  sipush  13107
15:  26:  invokevirtual#4; //Method printHelloWorld:(I)V
29:  29:  aload_1
30:  30:  sipush  17476
33:  33:  invokevirtual#4; //Method printHelloWorld:(I)V
36:  36:  aload_1
20:  37:  sipush  21845
40:  40:  invokevirtual#4; //Method printHelloWorld:(I)V
43:  43:  aload_1
44:  44:  sipush  26214
47:  47:  invokevirtual#4; //Method printHelloWorld:(I)V
25:  50:  return

```

Listing 2

Some virtual machines, including versions of the JRockit Virtual Machine
30 from BEA will compile this bytecode using standard compilation techniques. Some
of these techniques are described in the books "Advanced Compiler Design and
Implementation" by Steven S. Muchnik; "Crafting a Compiler with C" by Charles N.
Fischer and Richard J. LeBlanc, Jr.; and "Compilers" by Alfred V. Aho, Ravi Sethi,
Jeffrey D. Ullman, all of which are incorporated herein by reference. The process
35 is typically to translate the bytecode first into a high level intermediate
representation (HIR); then to a medium level intermediate representation (MIR);
and then to a low level intermediate level representation (LIR).

```

40 0x100080bc0:   push   %rbx
0x100080bc1:   mov    $0x140c2c0,%rax
0x100080bcb:   callq  0x100005410
0x100080bd0:   mov    %rsi,%rbx
0x100080bd3:   callq  0x100080720
0x100080bd8:   mov    $0x1111,%eax
45 0x100080bdd:   mov    %rbx,%rsi
0x100080be0:   mov    (%rbx),%ecx

```

```

0x100080be3:    nop
0x100080be4:    callq 0x10008075e
0x100080be9:    mov    $0x2222,%eax
5  0x100080bee:    mov    %rbx,%rsi
0x100080bf1:    mov    (%rbx),%ecx
0x100080bf4:    nop
0x100080bf5:    callq 0x10008075e
0x100080bfa:    mov    $0x3333,%eax
10 0x100080bff:    mov    %rbx,%rsi
0x100080c02:    mov    (%rbx),%ecx
0x100080c05:    nop
0x100080c06:    callq 0x10008075e
0x100080c0b:    mov    $0x4444,%eax
15 0x100080c10:    mov    %rbx,%rsi
0x100080c13:    mov    (%rbx),%ecx
0x100080c16:    nop
0x100080c17:    callq 0x10008075e
0x100080c1c:    mov    $0x5555,%eax
20 0x100080c21:    mov    %rbx,%rsi
0x100080c24:    mov    (%rbx),%ecx
0x100080c27:    nop
0x100080c28:    callq 0x10008075e
0x100080c2d:    mov    %rbx,%rsi
25 0x100080c30:    mov    $0x6666,%eax
0x100080c35:    mov    (%rbx),%ecx
0x100080c38:    nop
0x100080c39:    callq 0x10008075e
0x100080c3e:    pop   %rbx
30 0x100080c3f:    retq

```

Listing 3

Listing 3 shows the native machine code generated by versions of JRockit from the example code for the x86_64/Linux platform. The address 0x10008075e is the address of the stub that will call the native c-function. The assembler code in Listing 3 follows the specification in Table 1 for the x86_64/Linux platform: the calling convention for normal Java to Java calls puts the object pointer inside rsi, and the first integer into eax. If a second pointer had been used, it would have been put into rdi, and a second integer would be put into edx. Also according to the Java calling convention, rbx is a callee save register which is why it is restored from the stack before the return.

```

JNIEXPORT void JNICALL
Java_HelloWorld_printHelloWorld (JNIEnv *env, jobject obj, jint x)
{
45  printf ("Hello World %d\n", x);
}

```

Listing 4

Listing 4 shows an example of a function that follows the JNI-specification (previously referenced above) and which can therefore be called using the above native call.

5 A traditional native call stub

To interface the calls in the compiled Java code in Listing 3 to the c-function in Listing 4 on the x86_64/Linux platform, the stub needs to put a pointer to the JNIEnv in edi, leave the object pointer in esi, and move the integer argument from eax to ecx, before calling Java_HelloWorld_printHelloWorld. The traditional stub generator is tailored for each CPU/OS platform and generates the lowest level of IR (LIR) that is essentially a one-to-one mapping to machine code.

```

Variables, constants and labels:
15  v1 (reg, rsi)    v2 (reg, rax)    v3 (reg, rsp)
    c4 (i64, 0x30)  v5 [rsp+0x28]    v6 (reg, rbx)
    v7 [rsp+0x20]  v8 (reg, rbp)    v9 [rsp]
    v10 (reg, r10) c11 (i64, 0x6ab140) v12 [rsp+0xffffffffffffff00]

20  v13 [rbp+0x8]    v14 [rbp+0x1f0]
    v15 [rsp+0x8]    v16 (reg, r12)
    c17 (i64, 0xf)   v18 [rbx]
    c19 (i64, 0x8)   v20 [rbp+0x1e8]
    v21 (reg, rdi)   v22 (reg, rdx)
25  L23 (0x2aaaae1746d8) <--Address of Java_HelloWorld_printHelloWorld
    v24 [rbp+0x10]   L25 (0x100000580)
    c26 (i64, 0x0)   c27 (i64, 0xf4)

Parameters:
30  rsi (this) rax

block0:
35  0      x86_sub      rsp 0x30 ->rsp (i64)
    1      x86_mov      rbx ->[rsp+0x28] (i64)
    2      x86_mov      rbp ->[rsp+0x20] (i64)
    3      x86_mov      r10 ->r10 (i64)
    4      x86_mov      r10 ->[rsp] (i64)
    5      x86_mov      rax ->[rsp+0xffffffffffffff00] (i64)
40  6      lir_thread_vm ->rbp (i64)
    7      x86_mov      rsp ->[rbp+0x8] (i64)
    8      x86_mov      [rbp+0x1f0] ->rbx (i64)
    9      x86_mov      rbx ->[rsp+0x8] (i64)
    10     x86_mov      rsp ->r12 (i64)
    11     x86_mov      r12 ->r10 (i64)
45  12     x86_and      r10 0xf ->r10 (i64)
    13     x86_sub      rsp r10 ->rsp (i64)
    14     x86_push     rax
    15     x86_test     rsi rsi (i64)

```

```

== (then block3, else block4)
block1:
  34          x86_cmp          [rbp+0x8] 0x0 (i64)
== (then block7, else block8)
5  block2:
  33          x86_call         L25
goto (block1)
block3:
  16          x86_push         rsi
10 goto (block5)
block4:
  17          x86_mov          rsi ->[rbx] (i64)
  18          x86_push         rbx
  19          x86_add          rbx 0x8 ->rbx (i64)
15 block5:
  20          x86_lea          [rbp+0x1e8] ->r10 (i64)
  21          x86_push         r10
  22          x86_pop          ->rdi (i64)
  23          x86_pop          ->rsi (i64)
20  24          x86_pop          ->rdx (i64)
  25          x86_mov          rbx ->[rbp+0x1f0] (i64)
  26          lir_clrreg       ->rbp (i64)
  27          x86_call         L23
  28          lir_thread_vm    ->rbp (i64)
25  29          x86_mov          r12 ->rsp (i64)
  30          x86_mov          [rbp+0x10] ->r10 (i64)
  31          x86_mov          r10 ->[rbp+0x8] (i64)
  32          x86_test         r10 r10 (i64)
!= (then block2, else block1)
30 block6:
  37          x86_mov          [rsp+0x8] ->r10 (i64)
  38          x86_mov          r10 ->[rbp+0x1f0] (i64)
  39          x86_mov          [rsp+0x28] ->rbx (i64)
  40          x86_mov          [rsp+0x20] ->rbp (i64)
35  41          x86_add          rsp 0x30 ->rsp (i64)
  42          x86_ret
block7:
  35          x86_cmp          [rbp+0x10] 0x0 (i64)
== (then block6, else block8)
40 block8:
  36          lir_code         0xf4
goto (block6)

```

Listing 5

Listing 5 shows the low-level intermediate representation of the generated stub. As can be seen, the traditional stub generator has to produce the correct native instructions for the current CPU/OS platform. Only the branch commands are automatically added by the translation from LIR to native machine code.

Pseudo Code For A Traditional Native Call Stub Generator

The following pseudo code describes a traditional stub generator that is

tailored for the x86_64/Linux platform. It makes use of the following types:

pd addr is a C-storage type for platform dependent addresses.
64 bit long on the EM64T/AMD64 platform.

- 5 NativeFrame is a structure which is placed on the stack as part of the native call. It contains:
- oldHandles: is a pointer to the old JNI handles
 - preserved[PLATFORM_NOOF_PRESERVED_STORAGES]: space for all registers that need to be preserved over the native call.
 - 10 retAddr: Pointer to the return address in the compiled java code.
 - debugInfo: is a pointer to the debugInfo,
- The native frame also contains copies of the parameter registers and the return value register.

15

Variables:

- jci: is the JNI call info structure used for bookkeeping during stub generation.
- 20 preserved: is a list of all preserved storages for the java calling convention.
- np: is the length of the preserved list.
- ir: is the tree of the intermediate representation of the stub.
- current_block: contains code to setup call.
- 25 transitblock: is code necessary for the last transit to native code.
- end_block: is code to handle the return from native code.

Subroutines:

- 30 SetupParametersForCall: Transforms the java parameters to native parameters.
- GetFrom: *Acquire source var for LIR operation.*
- PushReference: Generate code that pushes a reference storage on the stack.
- 35 PushNativeParameter: Generate code that pushes a primitive type on the stack.

Step (1) The stub generator first sets up the book keeping data structures:

- 40 Setup return type of native function call.
`jci->ret_type`
- Setup how many storages are needed for the return type, int=1 or long=2.
`jci->ret_type_storages`
- 45 Setup `ret_frame` to reference storage location(s) for the java return type.

```
jci->ret_frame[]
```

Setup ret_callee to reference storage location(s) for the native return type.

```
5 jci->ret_callee[]
```

Later align is setup to contain reference to the IR op that performs stack alignment.

```
10 jci->align
```

```
jci->types= CG_MALLOC(env, (2*mpiGetNoofArgs(jci->
    mpi)+2)*sizeof(JlcType));
```

```
15 jci->pushed = 0;
```

```
jci->handles = 0;
```

```
jci->storage = CG_MALLOC(env, (2*mpiGetNoofArgs(jci->
    mpi)+2)*sizeof(Storage));
```

```
jci->jniEnv = env->jniEnv;
```

```
20
```

Step (2) Initialize three code blocks:

```
block = CreateIRBlock(ir);
```

```
end_block = CreateIRBlock(ir);
```

```
transit_block = CreateIRBlock(ir);
```

```
25 block_done = CreateIRBlock(ir);
```

Step (3) Allocate native frame. The initial part of the stack frame for native function calls is always the same size:

```
30 AppendIRToBlock(block, IR_X86_SUB(ir,
    IR_X86_SP,
    IR_CONSTANT(NATIVEFRAME_SIZE*sizeof(pd_addr)),
    IR_X86_SP))
```

```
35
```

Step (4) Store each preserved register on the stack so they can be restored after the JNI call:

```
for (i=0; i<np; ++i)
{
40     AppendIRToBlock(block, IR_X86_MOV(ir,
        irGetStorageVar(ir, preserved[i]),
        IR_RELATIVE_ID(X86_SP,
            sizeof(void*)*(np-i-1)+offsetof(
45     preserved))))
        struct NativeFrame,
}
```

Step (5) To make stack walking easier store a pointer to the native function in the stack frame:

```

AppendIRToBlock(block, IR_X86_MOV(ir,
5         IR_CONSTANT(irGetDebugInfo(ir)),
        JAVA_TRANSIT_SCRATCH_REG))
AppendIRToBlock(block, IR_X86_MOV(ir,
        JAVA_TRANSIT_SCRATCH_REG,
10         IR_RELATIVE_ID(X86_SP, offsetof(
                                struct NativeFrame,
                                debugInfo))))

```

Step (6) Add a unconditional write (stackbang) that verifies that it is at least one page of stack left for the native function:

```

15 AppendIRToBlock(block, IR_X86_MOV_T(ir, IR_PD_ADDR,
        IR_X86_A,
        IR_RELATIVE_ID(X86_SP, -(int)(size_t)sysPageSize)));

```

20 Step (7) Load the thread local data pointer and store the current stack pointer (the last Java stack frame) in the thread local data:

```

AppendIRToBlock(block, IR_LIR_THREAD_VM(ir, JAVA_TRANSIT_THREAD_REG));
AppendIRToBlock(block, IR_X86_MOV_T(ir, IR_PD_ADDR,
25         IR_X86_SP,
        JAVA_TRANSIT_THREAD_REG_ID(VMT_LAST_JAVA_FRAME_OFFSET)));

```

Step (8) Save the current JNI handles to the thread local data:

```

30 AppendIRToBlock(block, IR_X86_MOV_T(ir, IR_PD_ADDR,
        JAVA_TRANSIT_THREAD_REG_ID(VMT_HANDLES_OFFSET),
        JAVA2C_HANDLE_REG));
AppendIRToBlock(block, IR_X86_MOV_T(ir, IR_PD_ADDR,
        JAVA2C_HANDLE_REG,
35         IR_RELATIVE_ID(X86_SP, offsetof(
                                struct NativeFrame,
                                oldHandles))))

```

40 Step (9) Store the value of rsp before the parameters and alignment are pushed:

```

AppendIRToBlock(block, IR_X86_MOV_T(ir, IR_PD_ADDR,
        IR_X86_SP,
        JAVA2C_SCRATCH_PRESERVED_REG));

```

```
AppendIRToBlock(block, IR_X86_MOV_T(ir, IR_PD_ADDR,
    JAVA2C_SCRATCH_PRESERVED_REG,
    JAVA_TRANSIT_SCRATCH_REG));
```

5

Step (10) Align stack since the native calling convention X86_64 Linux requires 16 byte aligned stack pointers when entering a native function:

```
AppendIRToBlock(block, IR_X86_AND(ir,
    JAVA_TRANSIT_SCRATCH_REG,
    IR_CONSTANT(STACK_ALIGNMENT-1),
    JAVA_TRANSIT_SCRATCH_REG))
```

10

15

```
AppendIRToBlock(block, jci->align = IR_X86_SUB(ir,
    IR_X86_SP,
    JAVA_TRANSIT_SCRATCH_REG,
    IR_X86_SP))
```

20

Step (11) Place parameters in the correct places, basically push all on stack, and then pop some/all into the correct regs:

```
SetupParametersForCall(block);
```

25

Step (12) Store JNI handle in thread local data if the native function takes such arguments that need handles (usually pointers to objects):

30

```
if (jci.handles != 0)
{
    AppendIRToBlock(block, IR_X86_MOV_T(ir, IR_PD_ADDR,
        JAVA2C_HANDLE_REG,
        JAVA_TRANSIT_THREAD_REG_ID(VMT_HANDLES_OFFSET)))
}
```

Step (13) Patch the alignment of the stack if needed:

35

40

```
if (jci.pushed & 1) {
    InsertIROpAfter(IR_X86_SUB(ir,
        IR_X86_SP,
        IR_CONSTANT(sizeof(pd_addr)),
        IR_X86_SP),
        jci->align);
}
```

Step (14) The stack is setup, so the call can be made to the native function:

```
AppendIRToBlock(block, IR_X86_CALL(ir, IR_LABEL(ir, code)));
```

5 Step (15) Restore the stack pointer, and at the same time skip the stack alignment:

```
AppendIRToBlock(block, IR_X86_MOV_T(ir, IR_PD_ADDR,
    JAVA2C_SCRATCH_PRESERVED_REG,
10 IR_X86_SP));
```

Step (16) The thread local variable oldframe is a copy of the last Java frame if the native code has thrown exceptions. Otherwise oldframe will be NULL:

```
15 AppendIRToBlock(block, IR_X86_MOV_T(ir, IR_PD_ADDR,
    JAVA_TRANSIT_THREAD_REG_ID(VMT_CHECKATTRANSITFRAME_OFFSET),
    JAVA_TRANSIT_SCRATCH_REG));
```

Step (17) Restore the last Java frame:

```
20 AppendIRToBlock(block, IR_X86_MOV_T(ir, IR_PD_ADDR,
    JAVA_TRANSIT_SCRATCH_REG,
    JAVA_TRANSIT_THREAD_REG_ID(VMT_LAST_JAVA_FRAME_OFFSET)));
```

25 Step (18) Test if oldframe was zero or not. If oldframe is non-zero, an exception has happened and TransitToJava must be executed, otherwise proceed to the end_block:

```
30 AppendIRToBlock(block, IR_X86_TEST(ir,
    JAVA_TRANSIT_SCRATCH_REG,
    JAVA_TRANSIT_SCRATCH_REG));

AppendIRToBlock(block, IRBB_JCC, IR_NE);
ConnectIRBlock(ir, block, transit_block);
35 ConnectIRBlock(ir, block, end_block);
```

Step (19) Create the transit block that executes the TransitToJava function that takes care of thrown JNI exceptions. First store the native return value in the frame so it will survive TransitToJava:

```

for (i=0; i<jci->ret_type_storages; i++)
{
    from = irGetStorageVar(ir, data.ret_callee[i]);
    to   = IR_RELATIVE_ID(X86_SP, offsetof(struct NativeFrame,
5 space[i]));

    switch (CG_GET_STORAGE_TYPE(data.ret_callee[i]))
    {
        case STORAGE_TYPE_NORMAL:
10 AppendIRToBlock(transit_block, IR_X86_MOV_T(ir, IR_PD_ADDR,
from, to));
        break;
        case STORAGE_TYPE_FLOAT:
15 AppendIRToBlock(transit_block, IR_X86_FSTP_T(ir, jci.ret_type
== JLC_FLOAT ? IR_F32 : IR_F64,
from, to));
        break;
        case STORAGE_TYPE_XMM:
20 AppendIRToBlock(transit_block, jci.ret_type == JLC_FLOAT ?
IR_X86_MOVSS(ir, from, to) :
IR_X86_MOVSD(ir, from, to));
        break;
    }
}
25 AppendIRToBlock(block, IR_X86_CALL(ir, IR_LABEL(ir,
CI_GET_CODE(cgGetMethodCI(transitToJava_V))));

for (i=0 ; i<data.ret_type_storages ; i++)
30 {
    from = IR_RELATIVE_ID(X86_SP, offsetof(struct NativeFrame,
space[i]));
    to   = irGetStorageVar(ir, data.ret_callee[i]);

35 switch (CG_GET_STORAGE_TYPE(data.ret_callee[i]))
    {
        case STORAGE_TYPE_NORMAL:
AppendIRToBlock(block, IR_X86_MOV_T(ir, IR_PD_ADDR,
40 from, to));
        break;
        case STORAGE_TYPE_FLOAT:
AppendIRToBlock(block, IR_X86_FLD_T(ir, data.ret_type ==
JLC_FLOAT ? IR_F32 : IR_F64, from, to));
        break;
45 case STORAGE_TYPE_XMM:
AppendIRToBlock(block,
data.ret_type == JLC_FLOAT ? IR_X86_MOVSS(ir, from,
to) : IR_X86_MOVSD(ir, from, to));
        break;
50 }
}
ConnectIRBlock(ir, transit_block, end_block);

```

55 Step (20) To translate the return value from the native code to Java code, assume that the variable is in the correct native return position and move it to the

correct Java return position with proper sign extension or zero extension. IF the jci->ret_type == JLC_VOID this step can be skipped:

```

from = irGetStorageVar(ir, jci->ret_callee[0]);
to = irGetStorageVar(ir, jci->ret_frame[0]);
5
switch (jci->ret_type)
{
    case JLC_BOOLEAN:
    case JLC_BYTE:
10     AppendIRToBlock(end_block, IR_X86_MOVSX_T(ir, IR_I8, from,
to));
        break;
    case JLC_SHORT:
15     AppendIRToBlock(end_block, IR_X86_MOVSX_T(ir, IR_I16, from,
to));
        break;
    case JLC_CHAR:
        AppendIRToBlock(end_block, IR_X86_MOVZX_T(ir, IR_UI16, from,
20     to));
        break;
    case JLC_OBJECT:
    case JLC_INTERFACE:
    case JLC_ARRAY:
25     AppendIRToBlock(block, IR_X86_MOV_T(ir, IR_REF, from, to));
        append_nullcheck(ir, to, end_block, &block_null,
&block_non_null);
        AppendIRToBlock(block_non_null, IR_X86_MOV_T(ir, IR_REF,
30     IR_MEM(ir, to, 0, IR_NONE, 0),
to));

        ConnectIRBlock(ir, block_non_null, block_null);
        *end_block = block_null;
        break;
    case JLC_FLOAT:
35     case JLC_DOUBLE:
        if (jci->ret_callee[0] == jci->ret_frame[0])
        {
            // Do nothing since native frame and java frame share
            // the same return convention.
40         }
        else
        {
            IRType type;
            //this must be x87->xmm or xmm->x87
45         tmp = IR_RELATIVE_ID(X86_SP,
offsetof(struct NativeFrame,
space[0]));
            type = jci->ret_type == JLC_FLOAT ? IR_F32 : IR_F64;
50         switch (CG_GET_STORAGE_TYPE(jci->ret_callee[0]))
        {
            case STORAGE_TYPE_FLOAT:
                AppendIRToBlock(*block,
55                 IR_X86_FSTP_T(ir, type, from, tmp));
                AppendIRToBlock(*block, type == IR_F32 ?
                    IR_X86_MOVSS(ir, tmp, to) :

```

```

IR_X86_MOVSD(ir, tmp, to));
break;
case STORAGE_TYPE_XMM:
5 AppendIRToBlock(*block, type == IR_F32 ?
IR_X86_MOVSS(ir, from, tmp) :
IR_X86_MOVSD(ir, from, tmp));
AppendIRToBlock(*block, IR_X86_FLD_T(ir, type,
10 tmp, to));
break;
}
}
break;
15 }

```

Step (21) Restore the old handle pointer from the native frame on the stack to the thread local data:

```

20 AppendIRToBlock(end_block, IR_X86_MOV_T(ir, IR_PD_ADDR,
IR_RELATIVE_ID(X86_SP,
offsetof(struct NativeFrame, oldHandles)),
JAVA_TRANSIT_SCRATCH_REG))
25 AppendIRToBlock(end_block, IR_X86_MOV_T(ir, IR_PD_ADDR,
JAVA_TRANSIT_SCRATCH_REG,
JAVA_TRANSIT_THREAD_REG_ID(VMT_HANDLES_OFFSET)))

```

Step (22) Restore the preserved registers:

```

30 for (i=0 ; i<np ; i++)
{
AppendIRToBlock(end_block, IR_X86_MOV_T(ir, IR_PD_ADDR,
IR_RELATIVE_ID(X86_SP, offsetof(struct NativeFrame,
preserved) +
35 sizeof(void*)*(np-i-1)),
irGetStorageVar(ir, preserved[i]]));
}

```

Step (23) Deallocate the native frame from the stack:

```

40 AppendIRToBlock(end_block, IR_X86_ADD(ir,
IR_X86_SP,
IR_CONSTANT(NATIVEFRAME_SIZE*sizeof(pd_addr)),
IR_X86_SP));
45

```

Step (24) Add the return instruction back to compiled Java code:

```

AppendIRToBlock(end_block, IR_X86_RET(ir));

```

The function shown below describes with pseudo code how the parameters are setup to follow the calling convention of the native function. The strategy is to walk through all parameters and push all parameters on the stack that are supposed to be on the stack following the calling convention. Then the parameters that should be in registers are also pushed on the stack but immediately afterwards popped into the correct registers.

```

SetupParametersForCall(IRBlockP *block)
{
10   Storage used[PLATFORM_NOOF_STORAGE_TYPES];
      MPIIterS iter;
      int i, n, pos=0, pos0;

      n = mpiGetNoofArgs(jci->mpi);
15   memset(used, 0, PLATFORM_NOOF_STORAGE_TYPES*sizeof(int));
      pos0 = 0;

      // Fetch the JNIEnvironment
      jci->storage[pos0] =
20   platformGetNativeParamStorage(STORAGE_TYPE_NORMAL,
                                   pos0, used);

      pos0++;
      for (mpiGetIterator(&iter, jci->mpi, MPIITER_STORAGES, FALSE);
           mpiIteratorHasMore(&iter);
           mpiIteratorNext(&iter))
25   {
          jci->storage[pos0] = platformGetNativeParamStorage(
              STORAGE_TYPE_FOR(iter.jlcType), pos0, used);
          pos0++;
30   }

      pos = pos0;
      for (mpiGetIterator(&iter, jci->mpi, MPIITER_STORAGES, TRUE);
           mpiIteratorHasMore(&iter);
           mpiIteratorNext(&iter))
35   {
          if (CG_IS_STACK_STORAGE(jci->storage[--pos]))
          {
40   PushNativeParameter(block, iter.jlcType, iter.storage);
          }
      }

      pos = pos0;
45   for (mpiGetIterator(&iter, jci->mpi, MPIITER_STORAGES, TRUE);
           mpiIteratorHasMore(&iter);
           mpiIteratorNext(&iter))
      {
50   if (!CG_IS_STACK_STORAGE(jci->storage[--pos])) {
          PushNativeParameter(block, iter.jlcType, iter.storage);
      }
      }
}

```

```

// Load the address of the JNIEnvironment
AppendIRToBlock(*block, IR_X86_LEA_T(ir, IR_PD_ADDR,
    JAVA_TRANSIT_THREAD_REG_ID(VMT_JNI_INTERFACE_OFFSET),
    JAVA_TRANSIT_SCRATCH_REG));
5
AppendIRToBlock(block, IR_X86_PUSH(ir, get_from(data, ir,
    irVarGetStorage(ir, JAVA_TRANSIT_SCRATCH_REG)));

10
jci->types[jci->pushed++] = JLC_INT;
memset(used, 0, PLATFORM_NOOF_STORAGE_TYPES*sizeof(int));
for (i=jci->pushed-1 ; i>=0 ; i--)
{
    JlcType      jlcType = jci->types[i];
15    Storage      storage = platformGetNativeParamStorage(
        STORAGE_TYPE_FOR(jlcType), i, used);
    StorageType t = CG_GET_STORAGE_TYPE(storage);
    IRVAR        var = IR_NONE;

20    // All stack storages are ensured to be at the end
    if (t == STORAGE_TYPE_STACK) {
        break;
    } else {
        if (t == STORAGE_TYPE_XMM) {
25            if (jlcType == JLC_FLOAT)
                {
                    // Defaults to 32 bit
                    AppendIRToBlock(*block,
                        IR_X86_MOVSS(ir, IR_RELATIVE_ID(X86_SP,0),
30                        irGetOutParamVar(ir, storage));
                }
            else if (jlcType == JLC_DOUBLE)
                {
                    // Defaults to 64 bit
35                    AppendIRToBlock(*block, IR_X86_MOVSD(ir,
                        IR_RELATIVE_ID(X86_SP,0),
                        irGetOutParamVar(ir, storage));
                }
            var = JAVA_TRANSIT_SCRATCH_REG;
40        }
        else if (t == STORAGE_TYPE_NORMAL)
            {
                var = irGetOutParamVar(ir, storage);
            }
45        AppendIRToBlock(*block, IR_X86_POP(ir, var));
        jci->pushed--;
    }
}
}
}

50
IRVAR GetFrom(Storage storage)
{
    if (CG_IS_STACK_STORAGE(storage))
    {
55        return IR_RELATIVE_ID(JAVA2C_STACKPARAM_STORAGE,
            sizeof(pd_addr) * (STACK_EXTRA(data) +
            CG_GET_STORAGE_INDEX(storage)));
    }
}

```

```

else
{
    return irGetStorageVar(ir, storage);
5 }

PushReference(IRBlockP *block, Storage storage)
{
    IRBlockP block_null, block_not_null, block_done;
    IRVAR from;
10  if (CG_IS_STACK_STORAGE(storage)) {
        AppendIRToBlock(*block, IR_X86_MOV_T(ir, IR_REF,
            GetFrom(data, ir, storage),
            JAVA_TRANSIT_SCRATCH_REG));
15  from = JAVA_TRANSIT_SCRATCH_REG;
    } else {
        from = get_from(data, ir, storage);
    }
    append_nullcheck(ir, from, *block, &block_null, &block_not_null);
20  block_done = CreateIRBlock(ir);
    irBBSetMustNotHaveSafepoint(block_done);
    AppendIRToBlock(block_null, IR_X86_PUSH(ir, from));
    irBBConnect(ir, block_null, block_done);
    PushReference(data, ir, &block_not_null, storage, FALSE);
25  irBBConnect(ir, block_not_null, block_done);
    *block = block_done;
}

PushNativeParameter(IRBlockP *block, JlcType type, Storage storage)
30 {
    if (JLCTYPE_IS_PRIMITIVE(type))
    {
        AppendIRToBlock(*block, IR_X86_PUSH(ir, GetFrom(data, ir,
35  storage)));
    }
    else
    {
        PushReference(block, storage);
    }
40  jci->types[jci->pushed++] = type;
}

```

Listing 6 shows the final native machine code for the stub. Each machine
45 code instruction maps to one LIR operation in Listing 5, except for the branch
instructions, the frame setup and exit that have been added by the compiler.

```

-- Entering code with object pointer in esi and integer argument in eax.
50 0x80c70: sub    $0x30,%rsp
    0x80c74: mov    %rbx,0x28(%rsp)
    0x80c79: mov    %rbp,0x20(%rsp)
    0x80c7e: mov    $0x6ab140,%r10
    0x80c88: mov    %r10,(%rsp)
    0x80c8c: mov    %rax,0xffffffffffff000(%rsp)
55 0x80c94: mov    %fs:0xd8,%rbp

```

```

0x80c9d:  mov    %rsp,0x8(%rbp)
0x80ca1:  mov    0x1f0(%rbp),%rbx
0x80ca8:  mov    %rbx,0x8(%rsp)
5  0x80cad:  mov    %rsp,%r12
0x80cb0:  mov    %r12,%r10
0x80cb3:  and    $0xf,%r10
0x80cb7:  sub    %r10,%rsp
-- Push the integer argument, to be tranferred into rdx.
0x80cba:  push   %rax
10 0x80cbb:  test   %rsi,%rsi
0x80cbe:  jne    0x80cc3
-- Push the object pointer, to be popped into the same register.
0x80cc0:  push   %rsi
0x80cc1:  jmp    0x80ccc
15 0x80cc3:  mov    %rsi,(%rbx)
0x80cc7:  push   %rbx
0x80cc8:  add    $0x8,%rbx
0x80ccc:  lea   0x1e8(%rbp),%r10
-- Push the calculated address to the JNIEnvironment.
20 0x80cd3:  push   %r10
-- Pop the JNIEnvironment into rdi
0x80cd5:  pop    %rdi
-- Pop the rsi into rsi.
0x80cd6:  pop    %rsi
25 -- Pop the integer argument into rdx.
0x80cd7:  pop    %rdx
0x80cd8:  mov    %rbx,0x1f0(%rbp)
0x80cdf:  xor    %rbp,%rbp
-- Call the native function using ip relative addressing.
30 0x80ce2:  callq  *74(%rip)          # 0x100080d32
0x80ce8:  mov    %fs:0xd8,%rbp
0x80cf1:  mov    %r12,%rsp
0x80cf4:  mov    0x10(%rbp),%r10
0x80cf8:  mov    %r10,0x8(%rbp)
35 0x80cfc:  test   %r10,%r10
0x80cff:  je     0x80d06
0x80d01:  callq  0x00580
0x80d06:  cmpq   $0x0,0x8(%rbp)
0x80d0b:  jne    0x80d2f
40 0x80d0d:  cmpq   $0x0,0x10(%rbp)
0x80d12:  jne    0x80d2f
0x80d14:  mov    0x8(%rsp),%r10
0x80d19:  mov    %r10,0x1f0(%rbp)
0x80d20:  mov    0x28(%rsp),%rbx
45 0x80d25:  mov    0x20(%rsp),%rbp
0x80d2a:  add    $0x30,%rsp
0x80d2e:  retq
0x80d2f:  hlt
50 0x80d30:  jmp    0x80d14

```

Listing 6

As can be seen from the pseudo code above, the traditional stub generator uses low-level knowledge of the current CPU/OS platform to generate a proper stub for calling native code from compiled byte code. The programmer encodes

this knowledge manually. Since the traditional stub generator generates low-level native machine code, no compiler optimization techniques can be applied to the generated stub. Therefore the current techniques used for translating calling conventions in VMs are unsatisfactory due to non-optimal performance for short native functions, and because of the large amount of manual work needed to both add new calling conventions and to maintain the existing set of calling conventions of a large range of CPU/OS combinations.

Summary of the Invention:

Disclosed herein is a system and method for generation of efficient calls from compiled byte code within a virtual machine (VM) to native code within or outside of the VM. In particular the invention can be used in a system that includes a Java Virtual Machine (JVM). As described above, current techniques used to interface dynamically compiled byte code with native libraries using stubs are unsatisfactory due to non-optimal performance and of the large amount of manual work needed to maintain the stub generation code. An embodiment of the present invention addresses this problem by providing a system and a method to generate MIR code for native call stubs that can be compiled into native machine code that uses the current CPU/OS platform calling convention.

In accordance with an embodiment, the invention provides a system for efficient native call generation in a virtual machine environment, comprising: a computer running an operating system with native libraries; a virtual machine for executing a software application; a memory space for the compiled byte code as part of the software application, a compiler with a calling convention adaptor that generates MIR code for a native calls that is later compiled into native machine code. In accordance with another embodiment, the calling convention adaptor can include the native call into the MIR of the caller function. In accordance with another embodiment, the calling convention adaptor can generate MIR code for native calls of several different kinds of calling conventions.

Brief Description of the Figures:

Figure 1 shows an illustration of a system in accordance with an embodiment of the present invention, that includes a compiler with a calling convention adaptor that generates efficient native code calls from virtual machine byte codes.

Figure 2 shows a schematic of the MIR code for a native call, generated by the calling convention adaptor.

Figure 3 shows a schematic of the ParamInfo data structure that is used by the compiler to compile the MIR code for a native call into native machine code.

Detailed Description:

Disclosed herein is a system and method for generation of efficient calls from compiled byte code within a virtual machine to native code within or outside of the virtual machine. Traditionally the JIT defines its own calling convention for compiled Java to Java calls. Java to Native calls are handled with platform specific stubs. A large amount of manual work is required to implement the code that generates these stubs since they are different for each OS/CPU combination.

Figure 1 shows an illustration of a system in accordance with an embodiment of the present invention, that includes an byte code compiler with a calling convention adaptor which controls the compilation process. As shown in Figure 1, the system 101 may include a computer or physical server 102, that includes a hardware 103, operating system 104, libraries for operating system and user interface access 105 and run-time environment or virtual machine 106. The virtual machine may be a Java Virtual Machine (JVM), examples of which include the Sun Microsystems Inc. Java Virtual Machine, and the JRockit product from BEA Systems, Inc. The JVM contains the byte code to be execute 107, together with a byte code compiler 108. The present invention extends the byte code compiler with a calling convention adaptor 109 and ParamInfo structures 110. The JVM also includes, or has access to, a memory space 111, which in a Java environment comprises the native code generated by the compiler 108. The native code version of the Java Application 112, can make use of different libraries, like the Java Standard Library 113, the Java SWT Library 114 and a

Native Application Library 115.

The Java Standard Library is supplied by the JVM and the functionality that make use of operating system resources need to call a Native JVM Library 116. A native call 117 follows the JNI calling convention for the platform. The JNI calling convention is used for the call from the Java Standard Library to the Native JVM Library, for the Java SWT Library to the Native SWT Library and for the call from the Java Application to the Native Application Library. These calls (117) are transferred through a native call stub and the present invention improves both this stub and the process of generating it. The present invention can also include the native call into the MIR of the calling function to avoid a separate stub altogether.

A call 118 within the Java Application follows the Java-to-Java calling convention for the platform, which is different from the JNI calling convention. Since the Java Standard Library is supplied by the JVM, it is possible for a JVM implementation to add a separate (hidden) way of calling 119 the operating system resources using the platform calling convention, without going through an intermediate JNI-compatible library.

```

Variables and constants:
20 v1 (reg, rsp)          v2 (ref, v2)
   v3 (i32, v3)         v4 (i64, v4)
   v5 (i64, v5)         v6 (i64, v6)
   v7 (i64, v7)         v8 (i64, v8)
   v9 (i64, v9)         v10 (i64, v10)
25 v10 (i64, v10)        v11 (reg, [rsp+0x28])
   v12 (i64, v12)       v13 (i32, v13)
   v14 (i32, v14)       v15 (i64, v15)
   v16 (i64, v16)       v17 (i64, v17)
   v18 (i64, v18)       v19 (i64, v19)
30 c1 (i64, 8)           c2 (i64, 32)
   c3 (i64, 7208720)    c4 (i64, 560)
   c5 (i64, 24)         c6 (i64, 16)
   c7 (i64, 552)        c8 (ref, (nil))
   c9 (i64, 40)         c10 (i64, 0)
35 c11 (i64, 15)        c12 (i64, 204)

Labels: (L1 is printHelloWorld)
L1 (0x2aab2e4296d8, 0)   L2 ((nil), 0)

Parameters:
40 v2 (this) v3

block200:
   0 (i64) thread_vm    -> v4
goto (block201)
    
```

```

block201:
  1 (N/A) call      {type=method} {method=stackBang()V}
  2 (i64) add      rsp 8 -> v5
  3 (i64) mov      v5 -> v6
5   4 (i64) st      rsp 32 rsp
  5 (i64) mov      7208720 -> v7
  6 (i64) st      rsp 8 v7
  7 (i64) ld      v4 560 -> v8
  8 (i64) st      rsp 24 v8
10  9 (i64) lir_movlbl L2 -> v9
 10 (i64) st      rsp 16 v9
 12 (i64) add      v4 552 -> v10
 13 (ref) cmp     v2 (nil)
== (then block205, else block206)
15

block202:
 28 (N/A) call     v4 {type=method} {method=transitToJava(J)V}
goto (block203)
20

block203:
 29 (i64) ld      v4 8 -> v17
 30 (i64) ld      v4 16 -> v18
 31 (i64) or      v17 v18 -> v17
 32 (i64) cmp     v17 0
25 == (then block209, else block210)

block204:
 18 (i64) cast    v3 -> v14
 19 (i32) mov     v14 -> v13
30  20 (N/A) call    v4 {type=method}
    {method=transitFromJava_pd(J)V}
 21 (i64) and     rsp 15 -> v15
 22 (i64) cmp     v15 0
35 == (then block207, else block208)

block205:
 16 (i64) mov     0 -> v12
 17 (ref) st      rsp 40 v12
40 goto (block204)

block206:
 14 (ref) st      rsp 40 v2
 15 (i64) add     rsp 40 -> v12
45 goto (block204)

block207:
 24 (i64) st      v4 8 v5
 11 (i32) call    v10 v12 v13 L1 {keepalive=[rsp+0x28]}
    {type=method}
50 {method=printHelloWorld(I)V}
 25 (i64) ld      v4 16 -> v16
 26 (i64) st      v4 8 v16
 27 (i64) cmp     v16 0
55 != (then block202, else block203)

block208:
 23 (N/A) lir_code 204

```

```

block209:
  34 (N/A) call          v4 {type=method}
    {method=transitToJava_pd(J)V}
  36 (i64) ld           rsp 24 -> v19
5   37 (i64) st           v4 560 v19
    goto (block2011)

block210:
10  33 (N/A) lir_code     204

block211:
  35 (N/A) return
  36

```

Listing 7

15

Figure 2 shows a schematic of the MIR code for a native call generated by the calling convention adaptor in accordance with an embodiment of the present invention. The exact contents of each block can be seen in listing 7. In block 200, a pointer (v4) to the thread structure is acquired. In block 201, a stackbang is added to guarantee a minimum amount of stack space for the native call. The stack frame is setup for the native call, allocate JNI-handles in the thread structure, a pointer to a JNIEnv is extracted (v10) and a null pointer check is added. Block 202 transfers the control flow to the Java-exception handler if an exception was thrown from the native code. Block 203 contains useful sanity checks used when debugging. Block 204 translates 32-bit parameter (v3) into a 64-bit native call parameter (v12) and calls a platform dependent leave Java-function which is a no-op on the x86_64/Linux platform. It also checks that the stack is properly aligned. Block 205 moves a null this-pointer into the JNI handle. Block 206 moves a non-null this-pointer into the JNI handle. Block 207 calls the native function using v10, v12 and v13 as arguments. If an exception was thrown from the native code, the control flow will continue to 202, otherwise to 203. Block 208 contains the debug trap that follows a failed stack alignment check (204). Block 209 calls a platform dependent return to Java-function, which is a no-op on the x86_64/Linux platform and frees the JNI-handles allocated in block 201. Block 210 contains the debug trap that follows a failed sanity check (203). Block 211 is the return operation for the stub.

```

Storage gregs[] = {X86_64_RDI, X86_64_RSI, X86_64_RDX,
                  X86_64_RCX, X86_64_R8, X86_64_R9};
Storage fregs[] = {X86_XMM(0), X86_XMM(1), X86_XMM(2),
                  X86_XMM(3), X86_XMM(4), X86_XMM(5),
5                      X86_XMM(6), X86_XMM(7)};

Bool piSetupStoragesNative(ParamInfoP pinfo)
{
10     int curgreg;
    int curfreg;
    int offset;
    int sindex;
    int i, n;

15     n      = piGetNoofArgs(pinfo);
    offset  = 1;
    sindex  = 0;
    curgreg = curfreg = 0;

20     for (i = 0; i < n; i++) {
        switch(piGetJlcType(pinfo, i)) {
            case JLC_FLOAT:
            case JLC_DOUBLE:
                if (curfreg < NOOF_FREGS)
25                     PI_SET_STORAGE(pinfo, sindex++, fregs[curfreg++]);
                else
                    PI_SET_STORAGE(pinfo, sindex++, X86_STACK(offset++));
                break;
            case JLC_LONG:
            case JLC_INT:
            case JLC_SHORT:
            case JLC_BYTE:
            case JLC_CHAR:
            case JLC_BOOLEAN:
            case JLC_OBJECT:
            case JLC_ARRAY:
35                 if (curgreg < NOOF_GREGS)
                    PI_SET_STORAGE(pinfo, sindex++, gregs[curgreg++]);
                else
40                     PI_SET_STORAGE(pinfo, sindex++, X86_STACK(offset++));
                break;
            default:
                assert(0);
                return FALSE;
45         }
    }

    // Ok, now the return type. should not be counted into
    // 'nStorages', even though we store it.
50     switch (piGetReturnType(pinfo)) {
        case JLC_VOID:
            break;
        case JLC_LONG:
        case JLC_INT:
55         case JLC_SHORT:
        case JLC_BYTE:
        case JLC_CHAR:
        case JLC_BOOLEAN:

```

```

    case JLC_OBJECT:
    case JLC_ARRAY:
        PI_SET_STORAGE(pinfo, sindex, X86_64_RAX);
        break;
5   case JLC_FLOAT:
    case JLC_DOUBLE:
        // Even if we don't use xmm for our java convention,
        // native always returns floats in xmm0. deal with it
10  PI_SET_STORAGE(pinfo, sindex, X86_XMM(0));
        break;
    case JLC_INTERFACE:
        assert(0);
        return FALSE;
15 }

    pinfo->nStorages = sindex;
    pinfo->nOnStack = offset;

20 }
    return TRUE;
}

```

Listing 8

Figure 3 shows a schematic of the ParamInfo data structure, in accordance with an embodiment of the present invention, which is used by the compiler to

 compile the MIR code for the native call into native machine code. When the

 compiler translates a medium level call representation to a low level call

 representation it will use the ParamInfo structure to add move operations to

 transfer the MIR call parameters into their proper position for the native call and to

 move the return value from the call to its proper variable. Listing 8 shows the

 details of the procedure piSetupStoragesNative which is used to setup the

 ParamInfo for the native calling convention on the x86_64/Linux platform. The

 ParamInfo structure 300 contains: the call type 301, which is one of:

 PARAMTYPE_JAVA, PARAMTYPE_NATIVE, PARAMTYPE_NATIVE_JNI; the

 number of arguments 302; the exact number of platform words 303 that are

 pushed on the stack to store the arguments that cannot be stored in registers; the

 exact number of MIR storages 304 for the arguments (the exact number varies

 with the OS/CPU combination); the type for each parameter and the return value

 305, where each type is one of: JLC_LONG, JLC_DOUBLE, JLC_INT,

 JLC_FLOAT, JLC_SHORT, JLC_CHAR, JLC_BYTE, JLC_BOOLEAN, JLC_VOID,

 40 JLC_OBJECT, JLC_INTERFACE, JLC_ARRAY; the MIR storages 306 for the

 parameters, where a storage can be a register or a stack position and the MIR

storages (nRets) for the return type. nRets is zero if the return type is void. If the call is virtual, the first storage is the "this" pointer, which is then followed by the actual parameters. If the return type is not void, storages and a parameter may well share the same storage.

5 In accordance with an embodiment, the function piSetupStoragesNative takes as argument a ParamInfo filled with the correct call type (calltype), number of parameters (nParams), the parameter types and the return type (paramType). It then iterates over function parameters and creates proper storages for each parameter according to the calling convention, here encoded in the variables
10 gregs and fregs. It also creates storages for the return type. More than one storage for a parameter or the return type can be required, e.g. for 64-bit primitive data types on 32-bit platforms.

 When the compiler translates "call v10 v12 v13" in block 207:11 (illustrated in Listing 7) from MIR to LIR, it will iterate over the parameters (v10, v12, v13) and
15 for each parameter look up the proper storage for the parameter in the ParamInfo structure for the function to be called. If the storage is a normal integer register or stack position, a single LIR move operation is inserted to move the variable to the proper storage. If the storage is a floating point register that might require special handling on the platform, a few more LIR operations are added to allocate the
20 required resources for that particular parameter.

Pseudo Code For Improved Stub Generator

 The following pseudo code describes an improved stub generator in accordance with an embodiment of the present invention. It makes use of the
25 following types:

pd_addr is a C-storage type for platform dependent addresses.
64 bit long on the **EM64T/AMD64** platform.
30 NativeFrame is a structure which is placed on the stack as part of the native call. It contains:
 callAddr: the address of the native function to be called.
 oldHandles: is a pointer to the old JNI handles.
 realFrame: pointer to the real frame
35 debugInfo: a pointer used for stack walking

Variables:

mpi: ParamInfo structure used by the present invention to compile the MIR code.

params, nargs: The parameters to the function.

5 current_block, block_start, block_end, block_transit, block_return: Blocks
call, call_transit, ret, op, movlbl_native: Labels for MIR operations
thread_var: Points to the local data for the current thread.
ljf: Pointer to the last java frame on the stack.

10

Step (1) Add a variable that contains the pointer to the current thread.

(Block 200):

block_start = CreateIRBlock(ir);
thread_var = IR_VARIABLE(ir, IR_PD_ADDR);
15 AppendIRToBlock(block_start, IR_THREAD_VM(ir, thread_var));
[Block 200:0]

Step (2) Create more blocks:

20 current_block = CreateIRBlock(ir);
block_transit = CreateIRBlock(ir);
block_end = CreateIRBlock(ir);
ConnectIRBlock(ir, block_start, current_block);

25

Step (3) Add stack bang (Block 201):

AppendIRToBlock(current_block, IR_CALL_T(ir, IRFLAG_CALL_FIXED,
IR_NOTYPE,
cgGetCodeMethod(stackBang_V))

30 [Block 201:1]

Step (4) Allocate space on the stack for the native frame:

35 CI_SET_NOOF_PARAMS_ON_STACK(irGetDebufInfo(ir),
getNoofOnStackArgs(native_mpi));

nativeframe_pos = (FIRST_STACK_SLOT +
getNoofOnStackArgs(native_mpi)) * sizeof(void*);
40 handlestack_pos = nativeframe_pos + sizeof(struct NativeFrame);

Step (5) Calculate the pointer to the last Java frame, and fill it in. (Block 201):

```

ljf = IR_VARIABLE(ir, IR_PD_ADDR);
tmp = IR_VARIABLE(ir, IR_PD_ADDR);
5 [Block 201:2]
AppendIRToBlock(current_block, IR_ADD_T(ir, IR_PD_ADDR,
      sp,
      IR_CONSTANT(ir, IR_PD_ADDR, nativeframe_pos),
      ljf));
10 [Block 201:3]
AppendIRToBlock(current_block, IR_MOV_T(ir, IR_PD_ADDR,
      ljf,
      tmp));

```

15

Step (6) Store current stack pointer into the native frame on the stack. (Block 201):

```

AppendIRToBlock(current_block, STACK_STORE(ir, IR_PD_ADDR, sp,
20 nativeframe_pos +
      offsetof(struct NativeFrame, realFrame)));
[Block 201:4]

```

Step (7) Store the pointer to the debugInfo structure for the native function in the native frame as well. (Block 201):

```

tmp = IR_VARIABLE(ir, IR_PD_ADDR);
AppendIRToBlock(current_block, IR_MOV_T(ir, IR_PD_ADDR,
      IR_CONSTANT(ir, IR_PD_ADDR,
30 irGetDebugInfo(ir)),
      tmp));
[Block 201:5]
AppendIRToBlock(current_block, STACK_STORE(ir, IR_PD_ADDR,
      tmp,
      nativeframe_pos + offsetof(struct
35 NativeFrame,
      debugInfo)));
[Block 201:6]

```

Step (8) Store a pointer to the previous JNI handles in the native frame. (Block 201):

```

tmp = IR_VARIABLE(ir, IR_PD_ADDR);
AppendIRToBlock(current_block, IR_LD_T(ir, IR_PD_ADDR, IR_PD_ADDR,
      thread_var, IR_CONSTANT(ir, IR_PD_ADDR,

```



```

                                VMT_HANDLES_OFFSET), tmp));

```

```

[Block 201:7]

```

```

5 AppendIRToBlock(current_block, STACK_STORE(ir, IR_PD_ADDR,
                                tmp,
                                nativeframe_pos +
                                offsetof(struct NativeFrame,
                                        oldHandles)));

```

```

[Block 201:8]

```

```

10

```

Step (9) Store the address of the native function in the native frame, this is used for stack walking. (Block 201):

```

tmp = IR_VARIABLE(ir, IR_PD_ADDR);

```

```

15

```

```

AppendIRToBlock(current_block,
                movlbl_native = IR_LIR_MOVLBL(ir, IR_NONE, tmp));

```

```

[Block 201:9]

```

```

AppendIRToBlock(current_block,
                STACK_STORE(ir, IR_PD_ADDR,
                            tmp,
                            nativeframe_pos + offsetof(struct NativeFrame,
                                                        callAddr)));

```

```

20

```

```

[Block 201:10]

```

```

25

```

Step (10) Create the MIR operation for the call. Do not add it to any block yet.:

```

irGetIRTypeAndSizeForJlcType(piGetReturnType(mpi), &type, &s);
call = IR_CALL_T(ir, IRFLAG_CALL_FIXED, IR_SIZE2VARTYPE(type), method);
30 irOpAddSourceVar(ir, call, IR_LABEL(ir, getCallPtr(code)));

```

Step (11) Set the call to be of the Java native type:

```

IRCallSetInfoTypes(ir, call, CALLINFO_JAVA2C, PARAMTYPE_NATIVE_JNI);

```

```

35

```

Step (12) The first parameter to the native function is the JNI environment. (Block 201):

```

param = IR_VARIABLE(ir, IR_PD_ADDR);
40 AppendIRToBlock(current_block, IR_ADD_T(ir, IR_PD_ADDR, thread_var,
                                          IR_CONSTANT(ir, IR_PD_ADDR,
                                                      VMT_JNI_INTERFACE_OFFSET),
                                          param));

```

```
[Block 201:12]
IRAddSourceVar(ir, call, param);
```

- 5 Step (13) Add the parameters to the native function call. (Block 201,204,205,206):

```
for (pos=0 ; pos<nargs; pos++)
{
10       from = params[pos];
       type = irVarGetType(ir, from);

       if (type == IR_REF)
       {
15           block_done = CreateIRBlock(ir);
          append_nullcheck(ir, from, current_block, &block_null,
                            &block_non_null);
          [Block 201:13]

          ConnectIRBlock(ir, block_null, block_done);
20           ConnectIRBlock(ir, block_non_null, block_done);

          hsp = push_reference(ir, block_non_null, call, from,
                              &handlestack_pos, &param);
          [Block 206:14,15]
25           AppendIRToBlock(block_null, IR_MOV_T(ir, IR_PD_ADDR,
                                                  IR_0,
                                                  param));

          [Block 205:16,17]
          current_block = block_done;
30       }
       else
       {
          param = IR_VARIABLE(ir, type);
          if (type == IR_I32)
35           {
             tmp = IR_VARIABLE(ir, type);
             AppendIRToBlock(current_block, castOp =
                              SIGN_EXT_PARAMETER_TO_PLATFORM(ir, IR_PD_ADDR,
40                                                            from, tmp));
             [Block 204:18]
             from = tmp;
          }
          AppendIRToBlock(current_block, IR_MOV_T(ir, type, from, param));
45           [Block 204:19]
          }
          irOpAddSourceVar(ir, call, param);
       }
}
```

Step (14) Call the transit from Java function with the thread as an argument.
(Block 204, 208):

```

5 AppendIRToBlock(current_block, op = IR_CALL_T(ir, IRFLAG_CALL_FIXED,
  IR_NOTYPE,
                                cgGetCodeMethod(transitFromJava_pd_PV)));
  irOpAddSourceVar(ir, op, thread_var);
  [Block 204:20]
10 check_stack(ir, &current_block);
  [Block 204:21,22 8:23]

```

Step (15) The NativeFrame is now a Java-compatible stack frame, store this in the thread local area for garbage collection purposes. (Block 207):

```

15 AppendIRToBlock(current_block, IR_ST_T(ir, IR_PD_ADDR, IR_PD_ADDR,
  thread_var,
  IR_CONSTANT(ir, IR_PD_ADDR,
VMT_LAST_JAVA_FRAME_OFFSET),
  ljf));
20 [Block 207:24]

```

Step (16) Now add the call that was created earlier to the current block of code. (Block 207):

```

25 AppendIRToBlock(current_block, call);
  [Block 207:11]

```

Step (17) After the call, restore the previous Java frame. If the previous
30 Java frame was non-NULL it indicates that the process is returning to Java-code called from native JNI-code. Extra management is needed so transit2java is called. If an exception has occurred in the native code, then transit2java will need to be called. Otherwise, the function can return quickly. (Block 207):

```

35 frame_var = IR_VARIABLE(ir, IR_PD_ADDR);
  AppendIRToBlock(current_block, IR_LD_T(ir, IR_PD_ADDR, IR_PD_ADDR,
  thread_var,
  IR_CONSTANT(ir, IR_PD_ADDR,
VMT_CHECKATTRANSITFRAME_OFFSET),
  frame_var));
40 [Block 207:25]

```

```

AppendIRToBlock(current_block, IR_ST_T(ir, IR_PD_ADDR, IR_PD_ADDR,
    thread_var, IR_CONSTANT(ir, IR_PD_ADDR,
VMT_LAST_JAVA_FRAME_OFFSET),
    frame_var));
5 [Block 207:26]
AppendIRToBlock(current_block, IR_CMP_T(ir, IR_PD_ADDR, frame_var,
IR_0));
[Block 207:27]
irBBSetType(current_block, IRBB_JCC, IR_NE);
10 ConnectIRBlock(ir, current_block, block_transit);
ConnectIRBlock(ir, current_block, block_end);

```

Step (18) Call transit2java if an exception has happened. As can be seen
15 in this code the return object no longer needs to be explicitly preserved since the
normal register allocation and spilling takes care of protecting it for us. (Block
202):

```

current_block = block_transit;
AppendIRToBlock(transit_block,
20     call_transit = IR_CALL_T(ir, IRFLAG_CALL_FIXED, IR_NOTYPE,
        cgGetCodeMethod(transitToJava_PV)));
irOpAddSourceVar(ir, call_transit, thread_var);
[Block 202:28]
25 ConnectIRBlock(ir, block_transit, block_end);

```

Step (19) Make useful sanity checks. These can be ignored when not
debugging the JVM. (Block 203, 210):

```

30 check1 = IR_VARIABLE(ir, IR_PD_ADDR);
check2 = IR_VARIABLE(ir, IR_PD_ADDR);
current_block = block_end;
irBBAppendOp(current_block, IR_LD_T(ir, IR_PD_ADDR, IR_PD_ADDR,
    thread_var,
    IR_CONSTANT(ir, IR_PD_ADDR,
35 VMT_LAST_JAVA_FRAME_OFFSET),
    check1));
[Block 203:29]
irBBAppendOp(current_block, IR_LD_T(ir, IR_PD_ADDR, IR_PD_ADDR,
    thread_var,
40     IR_CONSTANT(ir, IR_PD_ADDR,
VMT_CHECKATTRANSITFRAME_OFFSET),
    check2));
[Block 203:30]
45 irBBAppendOp(current_block, IR_OR_T(ir, IR_PD_ADDR, check1, check2,
check1));
[Block 203:31]
append_nullcheck(ir, check1, block, &block_ok, &block_trap);
[Block 203:32]
irBBAppendOp(block_trap, IR_BREAK(ir));

```

```
[Block 210:33]
block_end = block_ok;
```

5 Step (20) Platform dependent transit2java part 2. This is a no-op on the x86_64/Linux platform. (Block 209):

```
current_block = block_end;
AppendIRToBlock(current_block, op = IR_CALL_T(ir, IRFLAG_CALL_FIXED,
10 IR_NOTYPE,
```

```
cgGetCodeMethod(transitToJava_pd_PV));
irOpAddSourceVar(ir, op, thread_var);
[Block 209:34]
```

15

Step (21) Calculate the return IR operation "ret" from the native function. In this example the return value is void:

```
if (piGetReturnType(mpi) == JLC_VOID)
20 {
    ret = IR_RETURN0(ir);
}
else
{
25     IRVAR res;
    irGetIRTypeAndSizeForJlcType(piGetReturnType(mpi), &flags, &s);

    type = IR_SIZE2VARTYPE(flags);
    res = IR_VARIABLE(ir, type);

30     if (type == IR_REF) {
        IRBlockP block_null, block_non_null;
        IRBlockP block_done = CreateIRBlock(ir);
        IRVAR    res_native = IR_VARIABLE(ir, IR_PD_ADDR);

35         irOpSetDestVar(call, res_native);
        append_nullcheck(ir, res_native, current_block, &block_null,
            &block_non_null);
        AppendIRToBlock(block_non_null, IR_LD_T(ir, IR_REF, IR_REF,
40 res_native,
            IR_0, res));
        AppendIRToBlock(block_null, IR_MOV_T(ir, IR_REF, IR_NULL, res));
        ConnectIRBlock(ir, block_non_null, block_done);
        ConnectIRBlock(ir, block_null, block_done);
        current_block = block_done;

45     } else {
        set_flags(ir, type);
        irOpSetDestVar(call, res);
    }
    ret = IR_RETURN_T(ir, type, res);
50 }
```

Step (22) Restore the JNI handles. (Block 209):

```

tmp = IR_VARIABLE(ir, IR_PD_ADDR);
AppendIRToBlock(current_block, STACK_LOAD(ir, IR_PD_ADDR,
5         nativeframe_pos + offsetof(struct NativeFrame,
         oldHandles),
         tmp));

[Block 209:36]
AppendIRToBlock(current_block, IR_ST_T(ir, IR_PD_ADDR, IR_PD_ADDR,
10         thread_var,
         IR_CONSTANT(ir, IR_PD_ADDR, VMT_HANDLES_OFFSET),
         tmp));

[Block 209:37]

```

15 Step (23) Now create the return code block. (Block 211):

```

block_return = CreateIRBlock(ir);
AppendIRToBlock(block_return, ret);
[Block 211:35]
20 ConnectIRBlock(ir, current_block, block_return);

```

Step (24) Now that the call IR operation exists the label can be inserted into the previous storage operation:

```

25 irOpSetSourceVar(movlbl_native, 0, irOpGetOrCreateLabel(ir, call));

-- Entering code with object pointer in rsi and integer argument in eax.
0x2aaaaee0e82d0:  push  %rbx
0x2aaaaee0e82d1:  push  %rbp
30 0x2aaaaee0e82d2:  sub   $0x38,%rsp
-- Save object pointer into r11.
0x2aaaaee0e82d6:  mov   %rsi,%r11
-- Save integer argument into esi
0x2aaaaee0e82d9:  mov   %eax,%esi
35 [Block 200]
0x2aaaaee0e82db:  mov   %fs:0xd8,%r13

[Block 201]
40 0x2aaaaee0e82e4:  movq  $0x0,0xffffffffffff000(%rsp)
0x2aaaaee0e82f0:  mov   %rsp,%r15
0x2aaaaee0e82f3:  add   $0x8,%r15
0x2aaaaee0e82f7:  mov   %r15,%r14
0x2aaaaee0e82fa:  mov   %rsp,0x20(%rsp)
0x2aaaaee0e82ff:  mov   $0x6dff10,%r15
45 0x2aaaaee0e8309:  mov   %r15,0x8(%rsp)
0x2aaaaee0e830e:  mov   0x230(%r13),%r8
0x2aaaaee0e8315:  mov   %r8,0x18(%rsp)
0x2aaaaee0e831a:  mov   $0x2aaaaee0e8371,%rax

```

```

0x2aaaae0e8324:  mov    %rax,0x10(%rsp)
0x2aaaae0e8329:  lea   0x228(%r13),%r15
0x2aaaae0e8330:  test  %r11,%r11
0x2aaaae0e8333:  jne   0x2aaaae0e833f
5
[Block 205]
0x2aaaae0e8335:  xor   %rdi,%rdi
0x2aaaae0e8338:  mov   %rdi,0x28(%rsp)
0x2aaaae0e833d:  jmp   0x2aaaae0e834b
10
[Block 206]
0x2aaaae0e833f:  mov   %r11,0x28(%rsp)
0x2aaaae0e8344:  mov   %rsp,%rdi
0x2aaaae0e8347:  add   $0x28,%rdi
15
[Block 204]
0x2aaaae0e834b:  mov   %r13,0x30(%rsp)
-- Move the integer argument from esi to rdx.
0x2aaaae0e8350:  movslq %esi,%rdx
20
0x2aaaae0e8353:  mov   %rsp,%r13
0x2aaaae0e8356:  and   $0xf,%r13
0x2aaaae0e835a:  test  %r13,%r13
0x2aaaae0e835d:  jne   0x2aaaae0e83c2
25
[Block 207]
0x2aaaae0e835f:  mov   0x30(%rsp),%rsi
0x2aaaae0e8364:  mov   %r14,0x8(%rsi)
-- Move calculated JNIEnv pointer into rsi.
0x2aaaae0e8368:  mov   %rdi,%rsi
30
-- Move object pointer into rdi.
0x2aaaae0e836b:  mov   %r15,%rdi
0x2aaaae0e836e:  mov   %rsp,%r12
0x2aaaae0e8371:  callq 0x2aab2e4296d8 <Java_HelloWorld_printHelloWorld>
0x2aaaae0e8376:  mov   %r12,%rsp
35
0x2aaaae0e8379:  mov   0x30(%rsp),%r10
0x2aaaae0e837e:  mov   0x10(%r10),%rcx
0x2aaaae0e8382:  mov   %rcx,0x8(%r10)
0x2aaaae0e8386:  test  %rcx,%rcx
0x2aaaae0e8389:  je    0x2aaaae0e839a
40
[Block 202]
0x2aaaae0e838b:  mov   0x30(%rsp),%rbx
0x2aaaae0e8390:  mov   %rbx,%rax
0x2aaaae0e8393:  callq 0x2aaaae0a2010
45
0x2aaaae0e8398:  jmp   0x2aaaae0e839f
[Block 203]
0x2aaaae0e839a:  mov   0x30(%rsp),%rbx
0x2aaaae0e839f:  mov   0x8(%rbx),%r9
50
0x2aaaae0e83a3:  mov   0x10(%rbx),%rdi
0x2aaaae0e83a7:  or    %rdi,%r9
0x2aaaae0e83aa:  test  %r9,%r9
0x2aaaae0e83ad:  jne   0x2aaaae0e83c3
55
[Block 209]
0x2aaaae0e83af:  mov   0x18(%rsp),%r12
0x2aaaae0e83b4:  mov   %r12,0x230(%rbx)

```

```
[Block 211]
0x2aaaaee0e83bb:  add    $0x38,%rsp
0x2aaaaee0e83bf:  pop    %rbp
0x2aaaaee0e83c0:  pop    %rbx
5  0x2aaaaee0e83c1:  retq

[Block 208]
0x2aaaaee0e83c2:  DEBUG trap
10 [Block 210]
0x2aaaaee0e83c3:  DEBUG trap
```

Listing 9

Listing 9 shows the final native machine code for the stub. In accordance with embodiments of the present invention, the virtual machine has enabled automatic generation of this stub. The stub translates the Java calling convention with the object pointer in rsi and the integer argument in eax to the native calling convention with a JNIEnv pointer in rsi, the object pointer in rdi and the integer argument in rdx. The calling convention adaptor creates the MIR code for this stub, not using platform specific knowledge. The compiler then compiles the MIR code using the ParamInfo which contained the platform specific knowledge. Since embodiments of the present invention extend the compiler to generate stubs (with the calling convention adaptor and the ParamInfo structure), there is no separate manual work required to maintain the stub generation code. If the compiler can compile for the platform, it can generate native call stubs for the platform. Furthermore, since the compiler has access to a higher level description of the native call stub it can apply advanced compiler optimization techniques on the stub to increase the code efficiency. It is even possible to include the MIR of the native call into the MIR of the caller to achieve even higher performance improvements. There is no limitation on what kind of native calls can be generated. JNI-compatible calls such as 117 in Figure 1 can co-exist with raw calls such as 119 in Figure 1 within the same JVM. Therefore embodiments of the present invention improve native call performance, and reduce the manual work needed to add new calling conventions when the number of CPU/OS combinations grow.

The present invention may be conveniently implemented using a conventional general purpose or a specialized digital computer or microprocessor programmed according to the teachings of the present disclosure. Appropriate software coding can readily be prepared by skilled programmers based on the

teachings of the present disclosure, as will be apparent to those skilled in the software art.

In some embodiments, the present invention includes a computer program product which is a storage medium (media) having instructions stored thereon/in
5 which can be used to program a computer to perform any of the processes of the present invention. The storage medium can include, but is not limited to, any type of disk including floppy disks, optical discs, DVD, CD-ROMs, microdrive, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, DRAMs, VRAMs, flash memory devices, magnetic or optical cards, nanosystems (including
10 molecular memory ICs), or any type of media or device suitable for storing instructions and/or data.

The foregoing description of the present invention has been provided for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and
15 variations will be apparent to the practitioner skilled in the art. Particularly, it will be evident that while the examples described herein illustrate how the features may be used in a WebLogic or JRockit environment, other application servers, virtual machines, JVMs, computing environments, and software development systems may use and benefit from the invention. The code examples given are
20 presented for purposes of illustration. It will be evident that the techniques described herein may be applied using other code languages, and with different code.

The embodiments were chosen and described in order to best explain the principles of the invention and its practical application, thereby enabling others
25 skilled in the art to understand the invention for various embodiments and with various modifications that are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the following claims and their equivalence.

Claims:

What is claimed is:

- 5 1. A system for generating native code calls from byte code in virtual machines, comprising:
- a computer running an operating system with native libraries;
 - a virtual machine for executing a software application;
 - a memory space for the compiled byte code as part of the software
- 10 application; and
- a compiler with a calling convention adaptor that integrates native calls into the high level analysis step of the compilation of byte codes.
2. The system of claim 1, wherein the compiler uses a ParamInfo structure to
- 15 compile a MIR call into a LIR call.
3. The system of claim 1, wherein the MIR of the native call is inserted into the MIR of the caller code.
- 20 4. The system of claim 1, wherein the MIR for Java to Java calls are compiled using the ParamInfo structure.
5. The system of claim 4, wherein several different native calling conventions
- 25 are used simultaneously.
6. A method for generating native code calls from byte code in virtual machines, comprising the steps of:
- providing a computer running an operating system with native libraries;
 - providing a virtual machine for executing a software application;
 - providing a memory space for the compiled byte code as part of the
- 30 software application; and

providing a compiler with a calling convention adaptor that integrates native calls into the high level analysis step of the compilation of byte codes.

7. The method of claim 6, wherein the compiler uses a ParamInfo structure to
5 compile a MIR call into a LIR call.
8. The method of claim 6, wherein the MIR of the native call is inserted into the MIR of the caller code.
- 10 9. The method of claim 6, wherein the MIR for Java to Java calls are compiled using the ParamInfo structure.
- 15 10. The method of claim 9, wherein several different native calling conventions are used simultaneously.

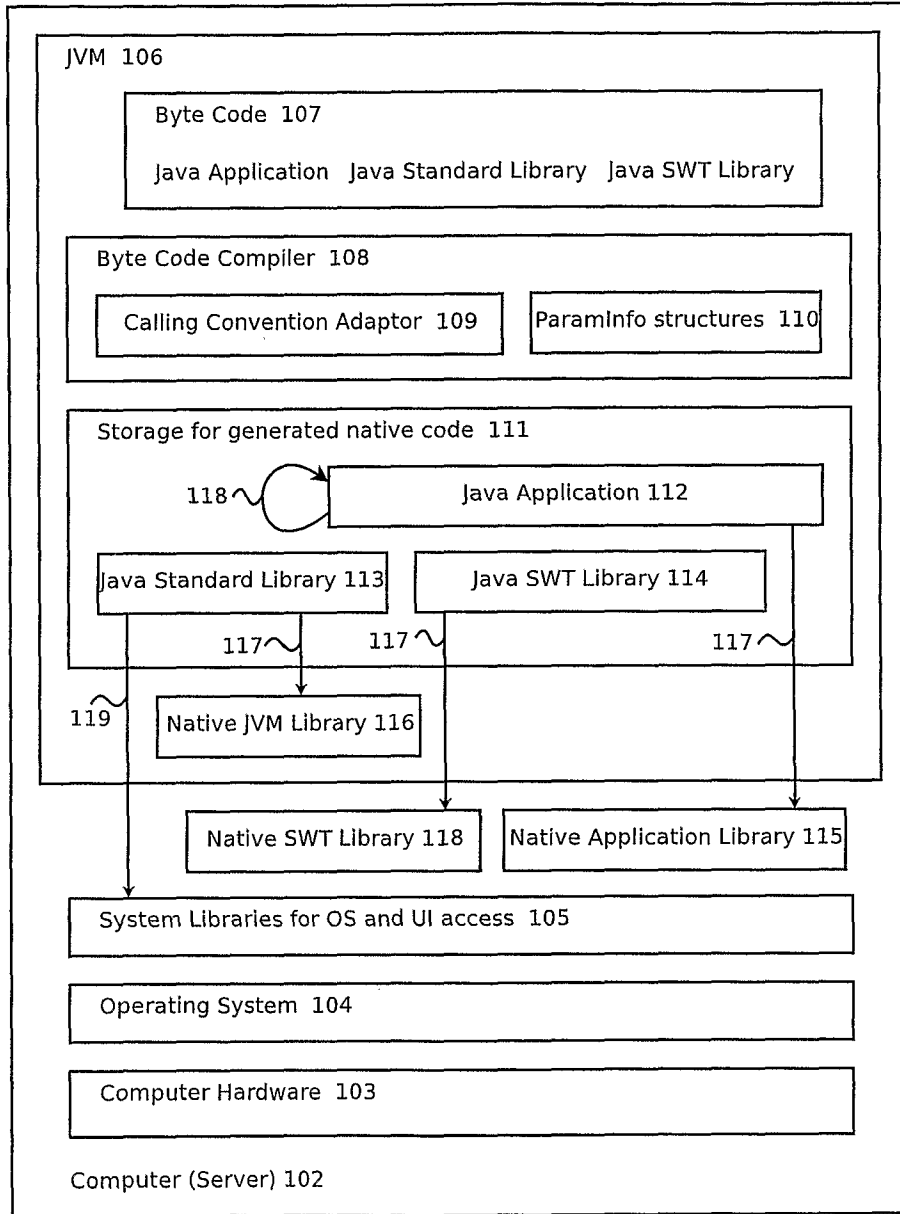


FIG. 1

101




FIG. 2

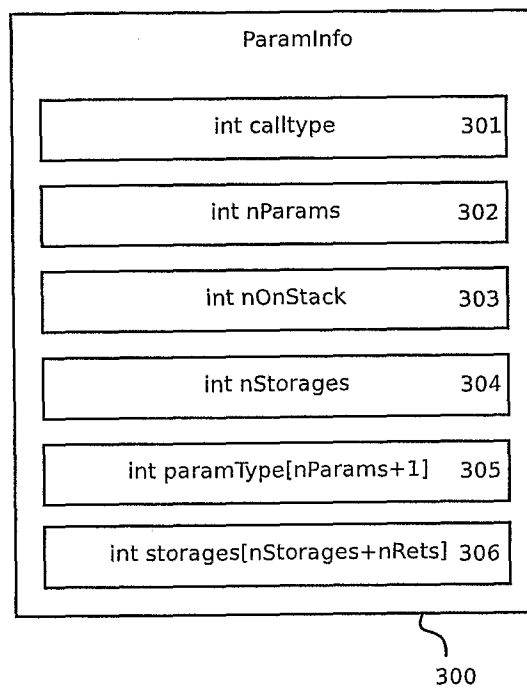


FIG. 3