



US 20070234070A1

(19) **United States**(12) **Patent Application Publication**
Horning et al.(10) **Pub. No.: US 2007/0234070 A1**(43) **Pub. Date: Oct. 4, 2007**(54) **SOFTWARE SELF-DEFENSE SYSTEMS AND METHODS****Publication Classification**(75) Inventors: **James J. Horning**, Palo Alto, CA (US);
W. Olin Sibert, Lexington, MA (US);
Robert E. Tarjan, Princeton, NJ (US);
Urnesh Maheshwari, San Jose, CA (US); **William G. Horne**,
Lawrenceville, NJ (US); **Andrew K. Wright**,
Monroe Township, NJ (US); **Lesley R. Matheson**,
Campbell, CA (US); **Susan Owicki**, Palo Alto, CA (US)(51) **Int. Cl.**
H04N 7/16 (2006.01)
G06F 12/14 (2006.01)
H04L 9/32 (2006.01)
G06F 11/30 (2006.01)
G06F 17/30 (2006.01)
G06F 7/04 (2006.01)
G06K 9/00 (2006.01)
H03M 1/68 (2006.01)
H04K 1/00 (2006.01)
H04L 9/00 (2006.01)(52) **U.S. Cl.** **713/190; 713/187; 726/26**

Correspondence Address:


**FINNEGAN, HENDERSON, FARABOW,
GARRETT & DUNNER
LLP****901 NEW YORK AVENUE, NW
WASHINGTON, DC 20001-4413 (US)**(73) Assignee: **Intertrust Technologies Corp.**, Sunnyvale, CA(21) Appl. No.: **11/506,321**(22) Filed: **Aug. 17, 2006****Related U.S. Application Data**

(63) Continuation of application No. 09/629,546, filed on Jul. 31, 2000.

(60) Provisional application No. 60/146,419, filed on Jul. 29, 1999.

(57) **ABSTRACT**

Systems and methods are disclosed for protecting a computer program from unauthorized analysis and modification. Obfuscation transformations can be applied to the computer program's local structure, control graph, and/or data structure to render the program more difficult to understand and/or modify. Tamper-resistance mechanisms can be incorporated into the computer program to detect attempts to tamper with the program's operation. Once an attempt to tamper with the computer program is detected, the computer program reports it to an external agent, ceases normal operation, and/or reverses any modifications made by the attempted tampering. The computer program can also be watermarked to facilitate identification of its owner. The obfuscation, tamper-resistance, and watermarking transformations can be applied to the computer program's source code, object code, or executable image.



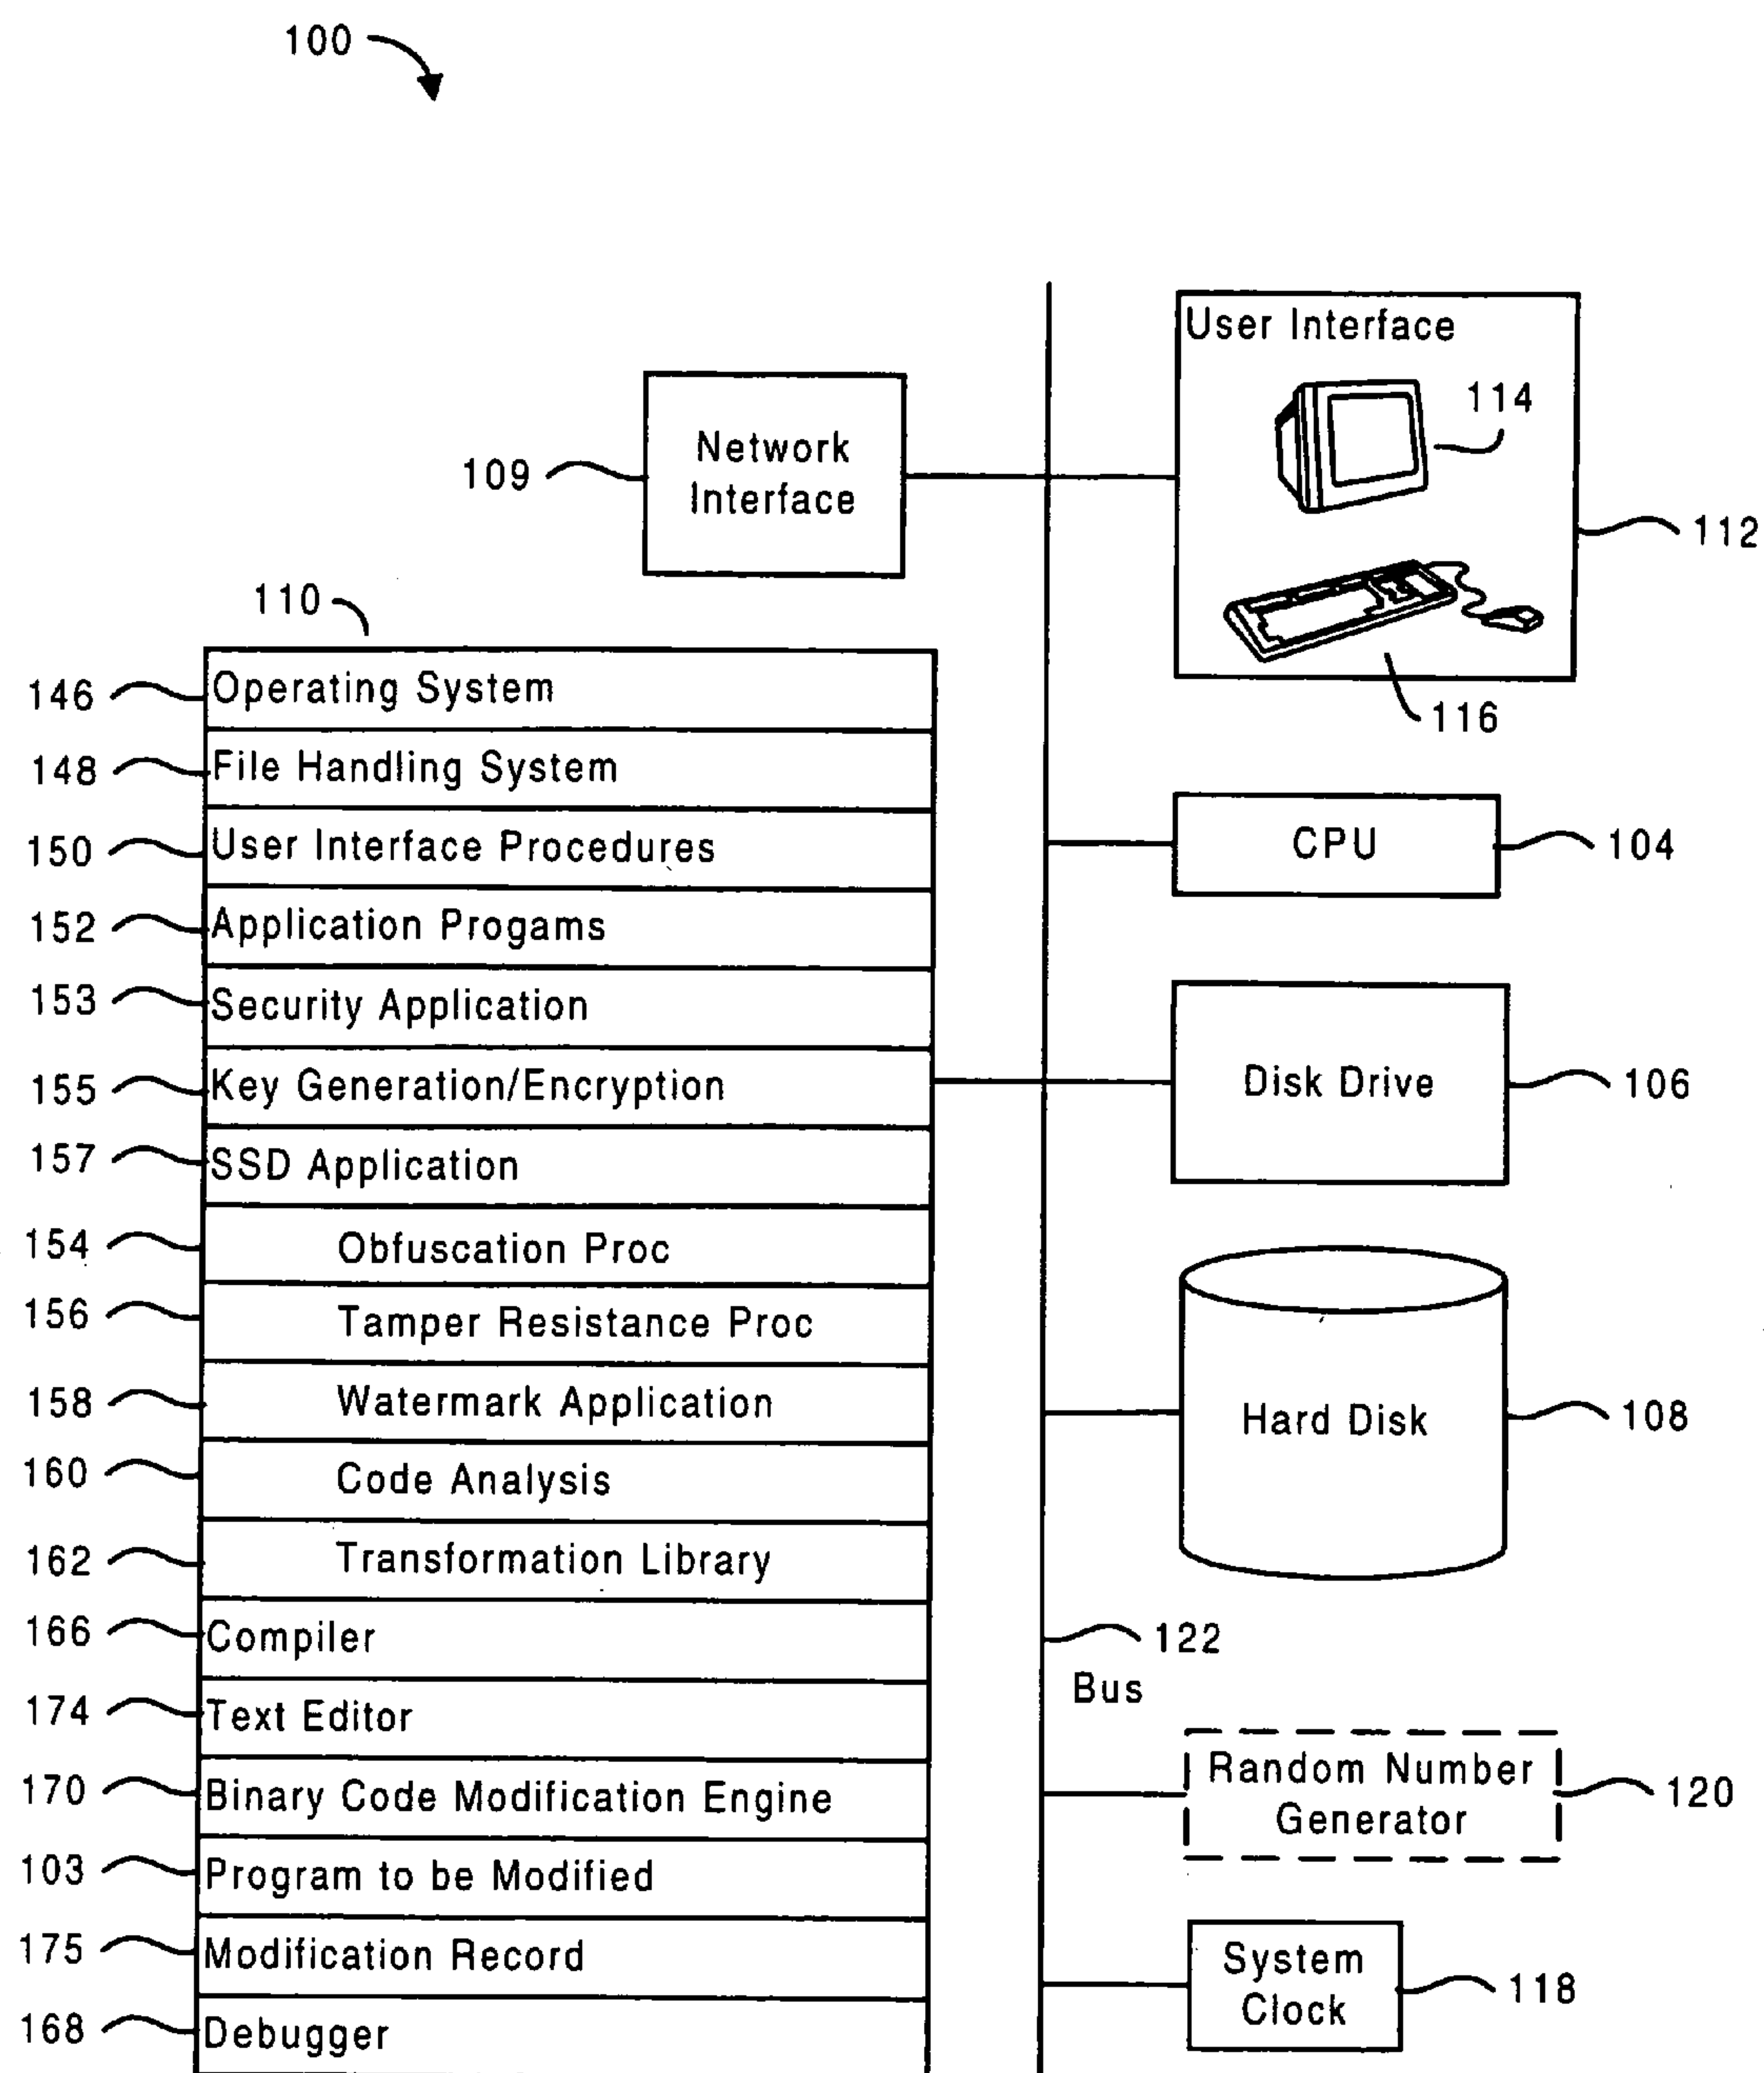


Fig. 1




FIG. 2




Fig. 3

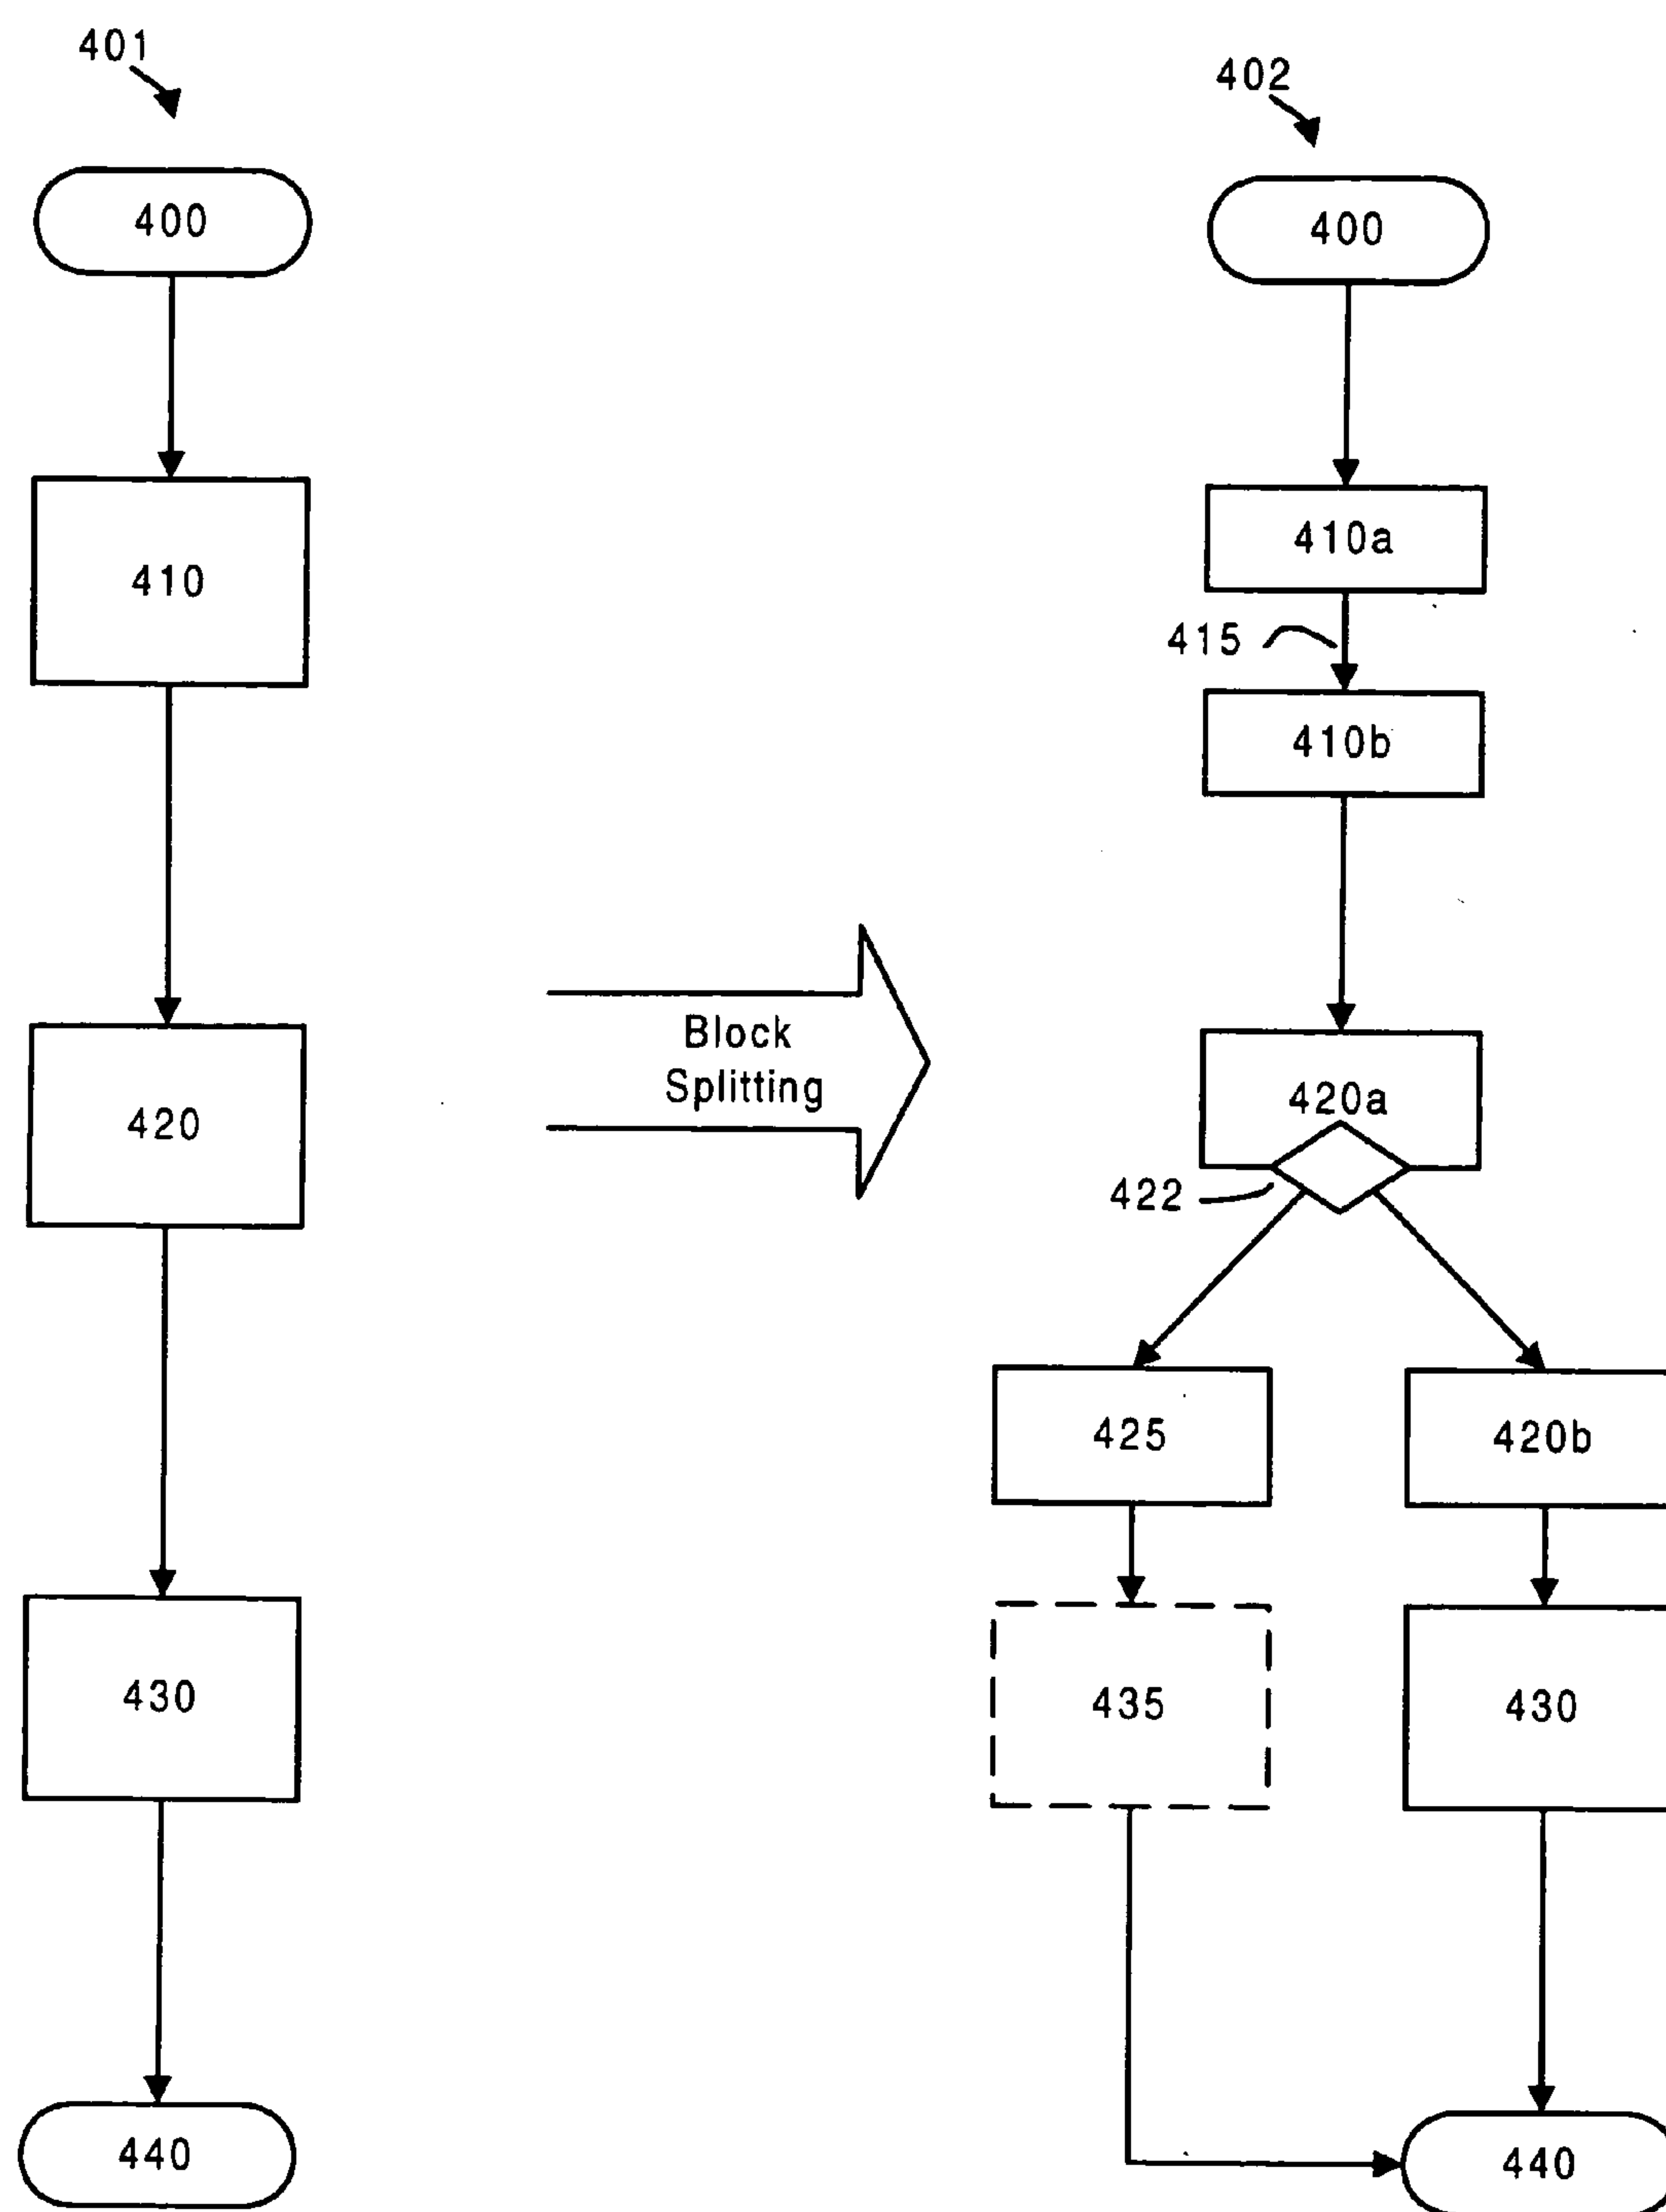


FIG. 4

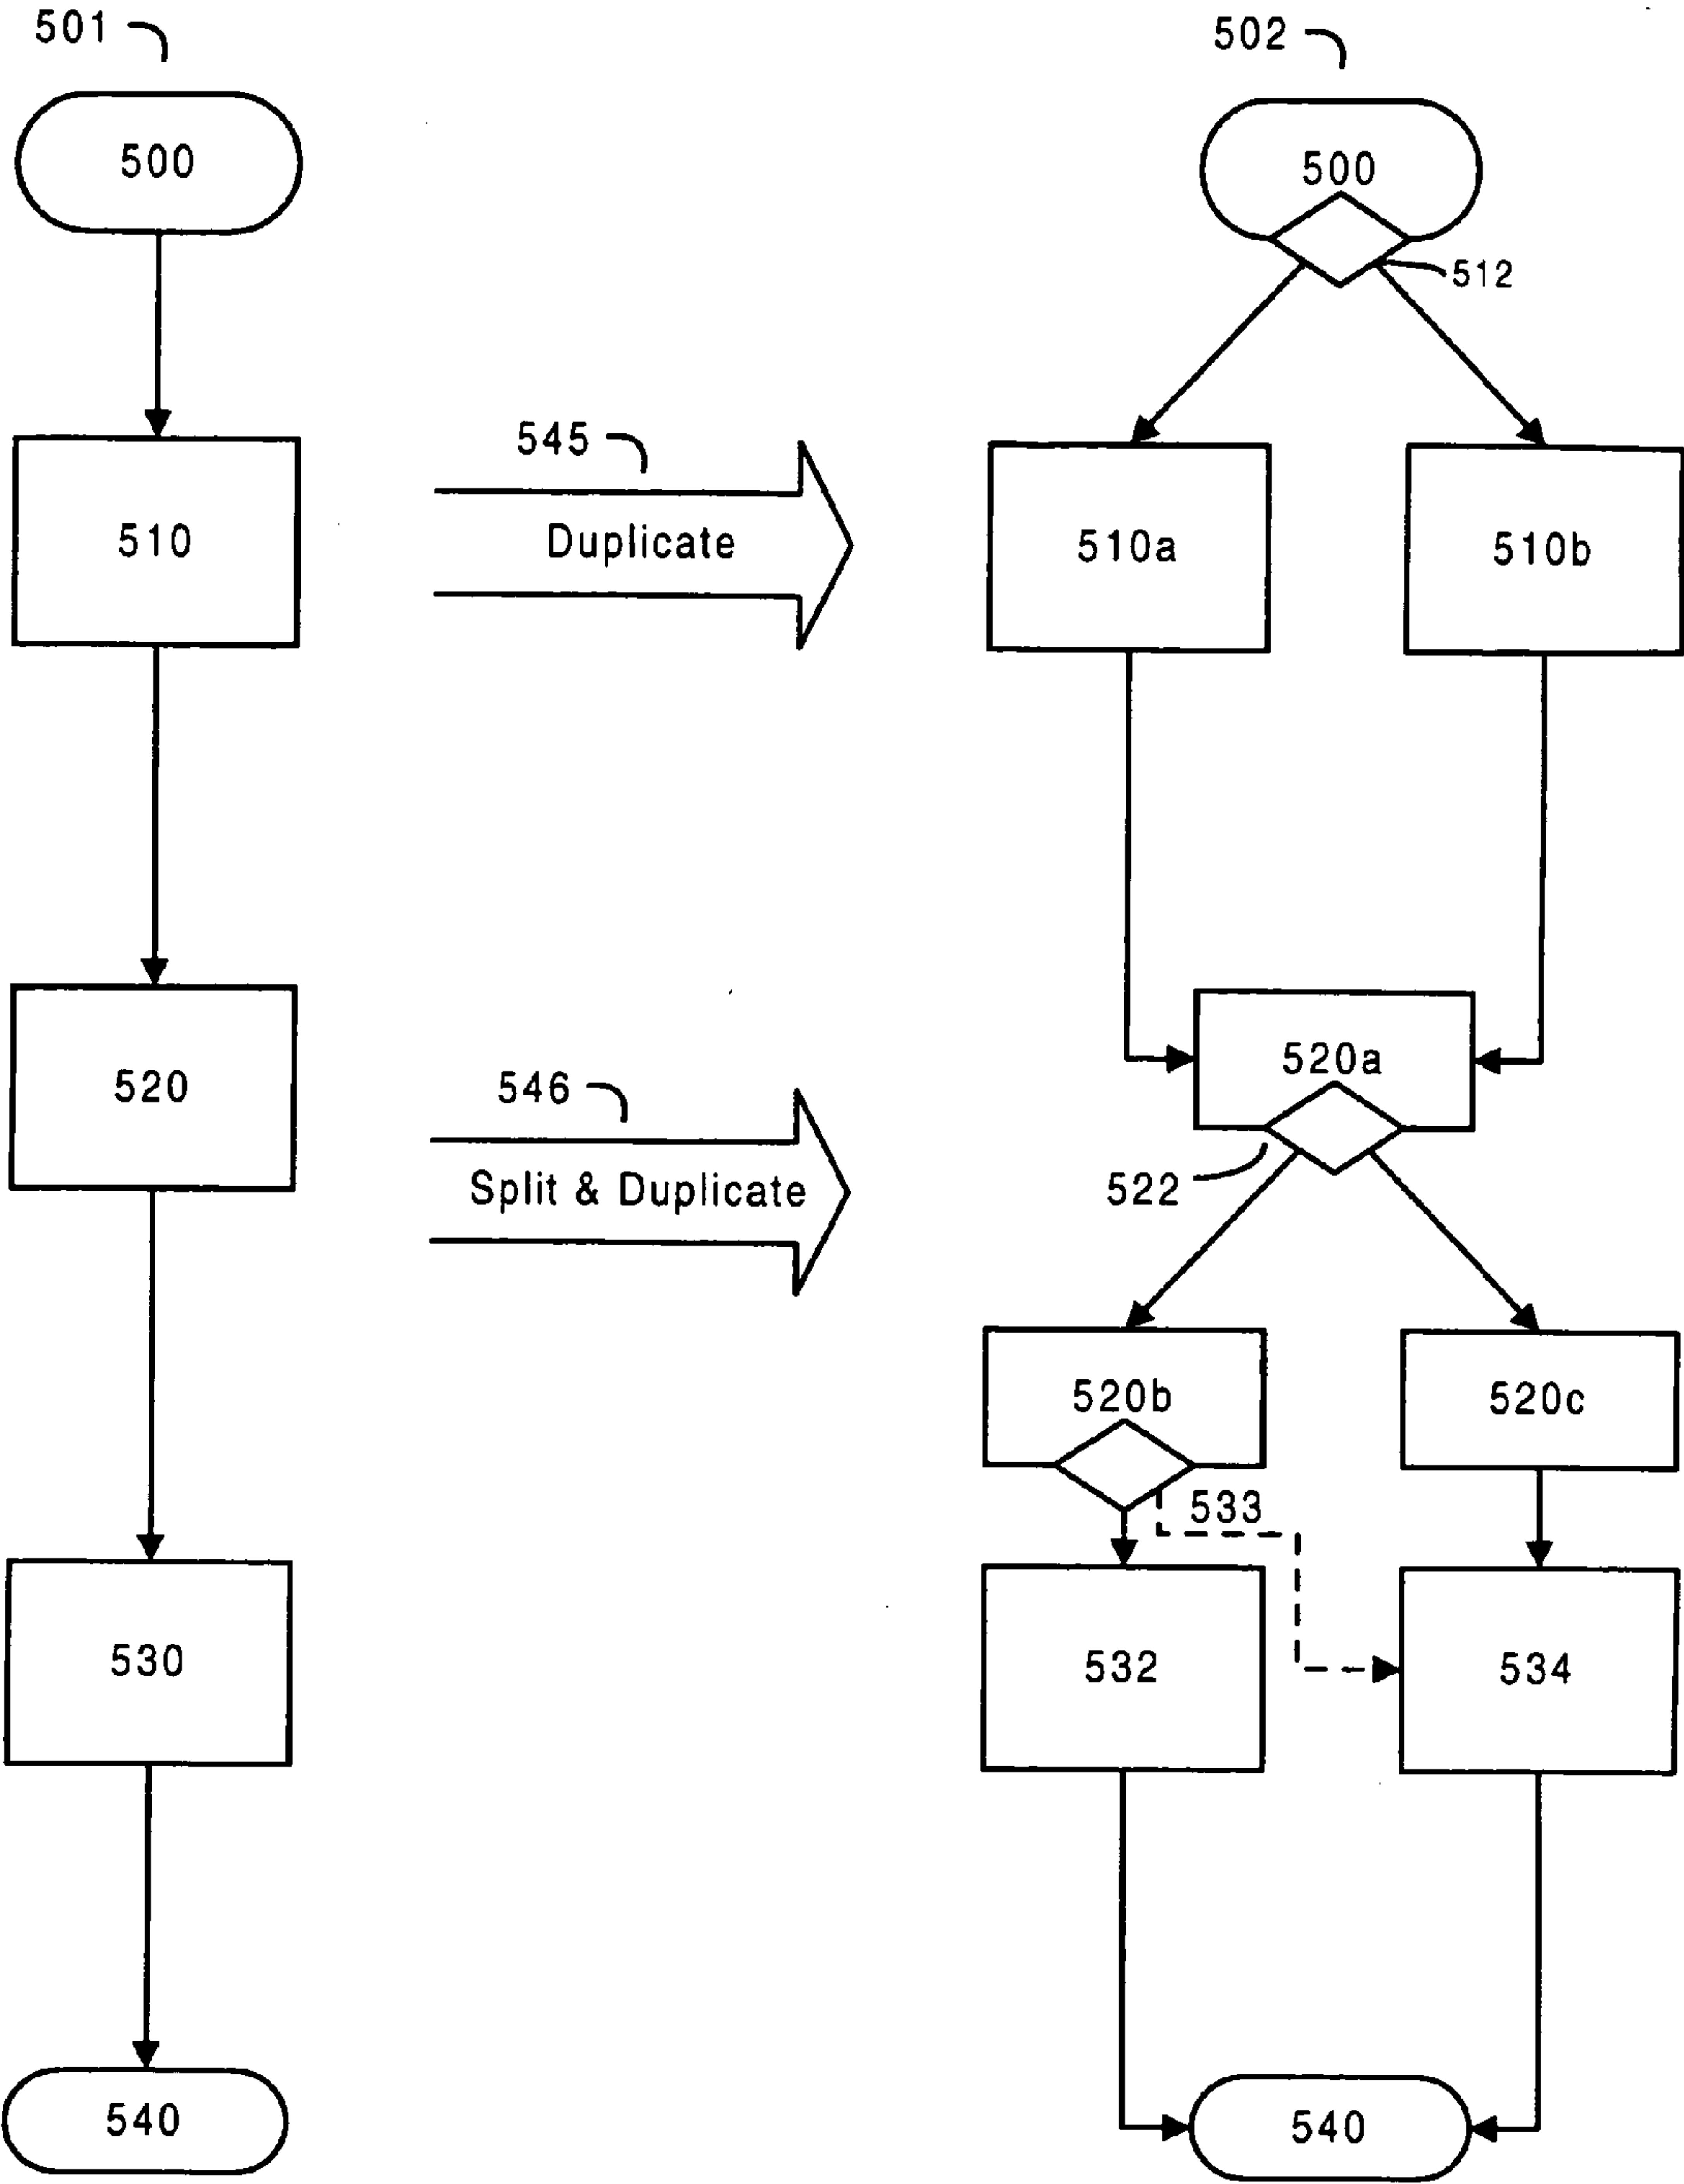


FIG. 5

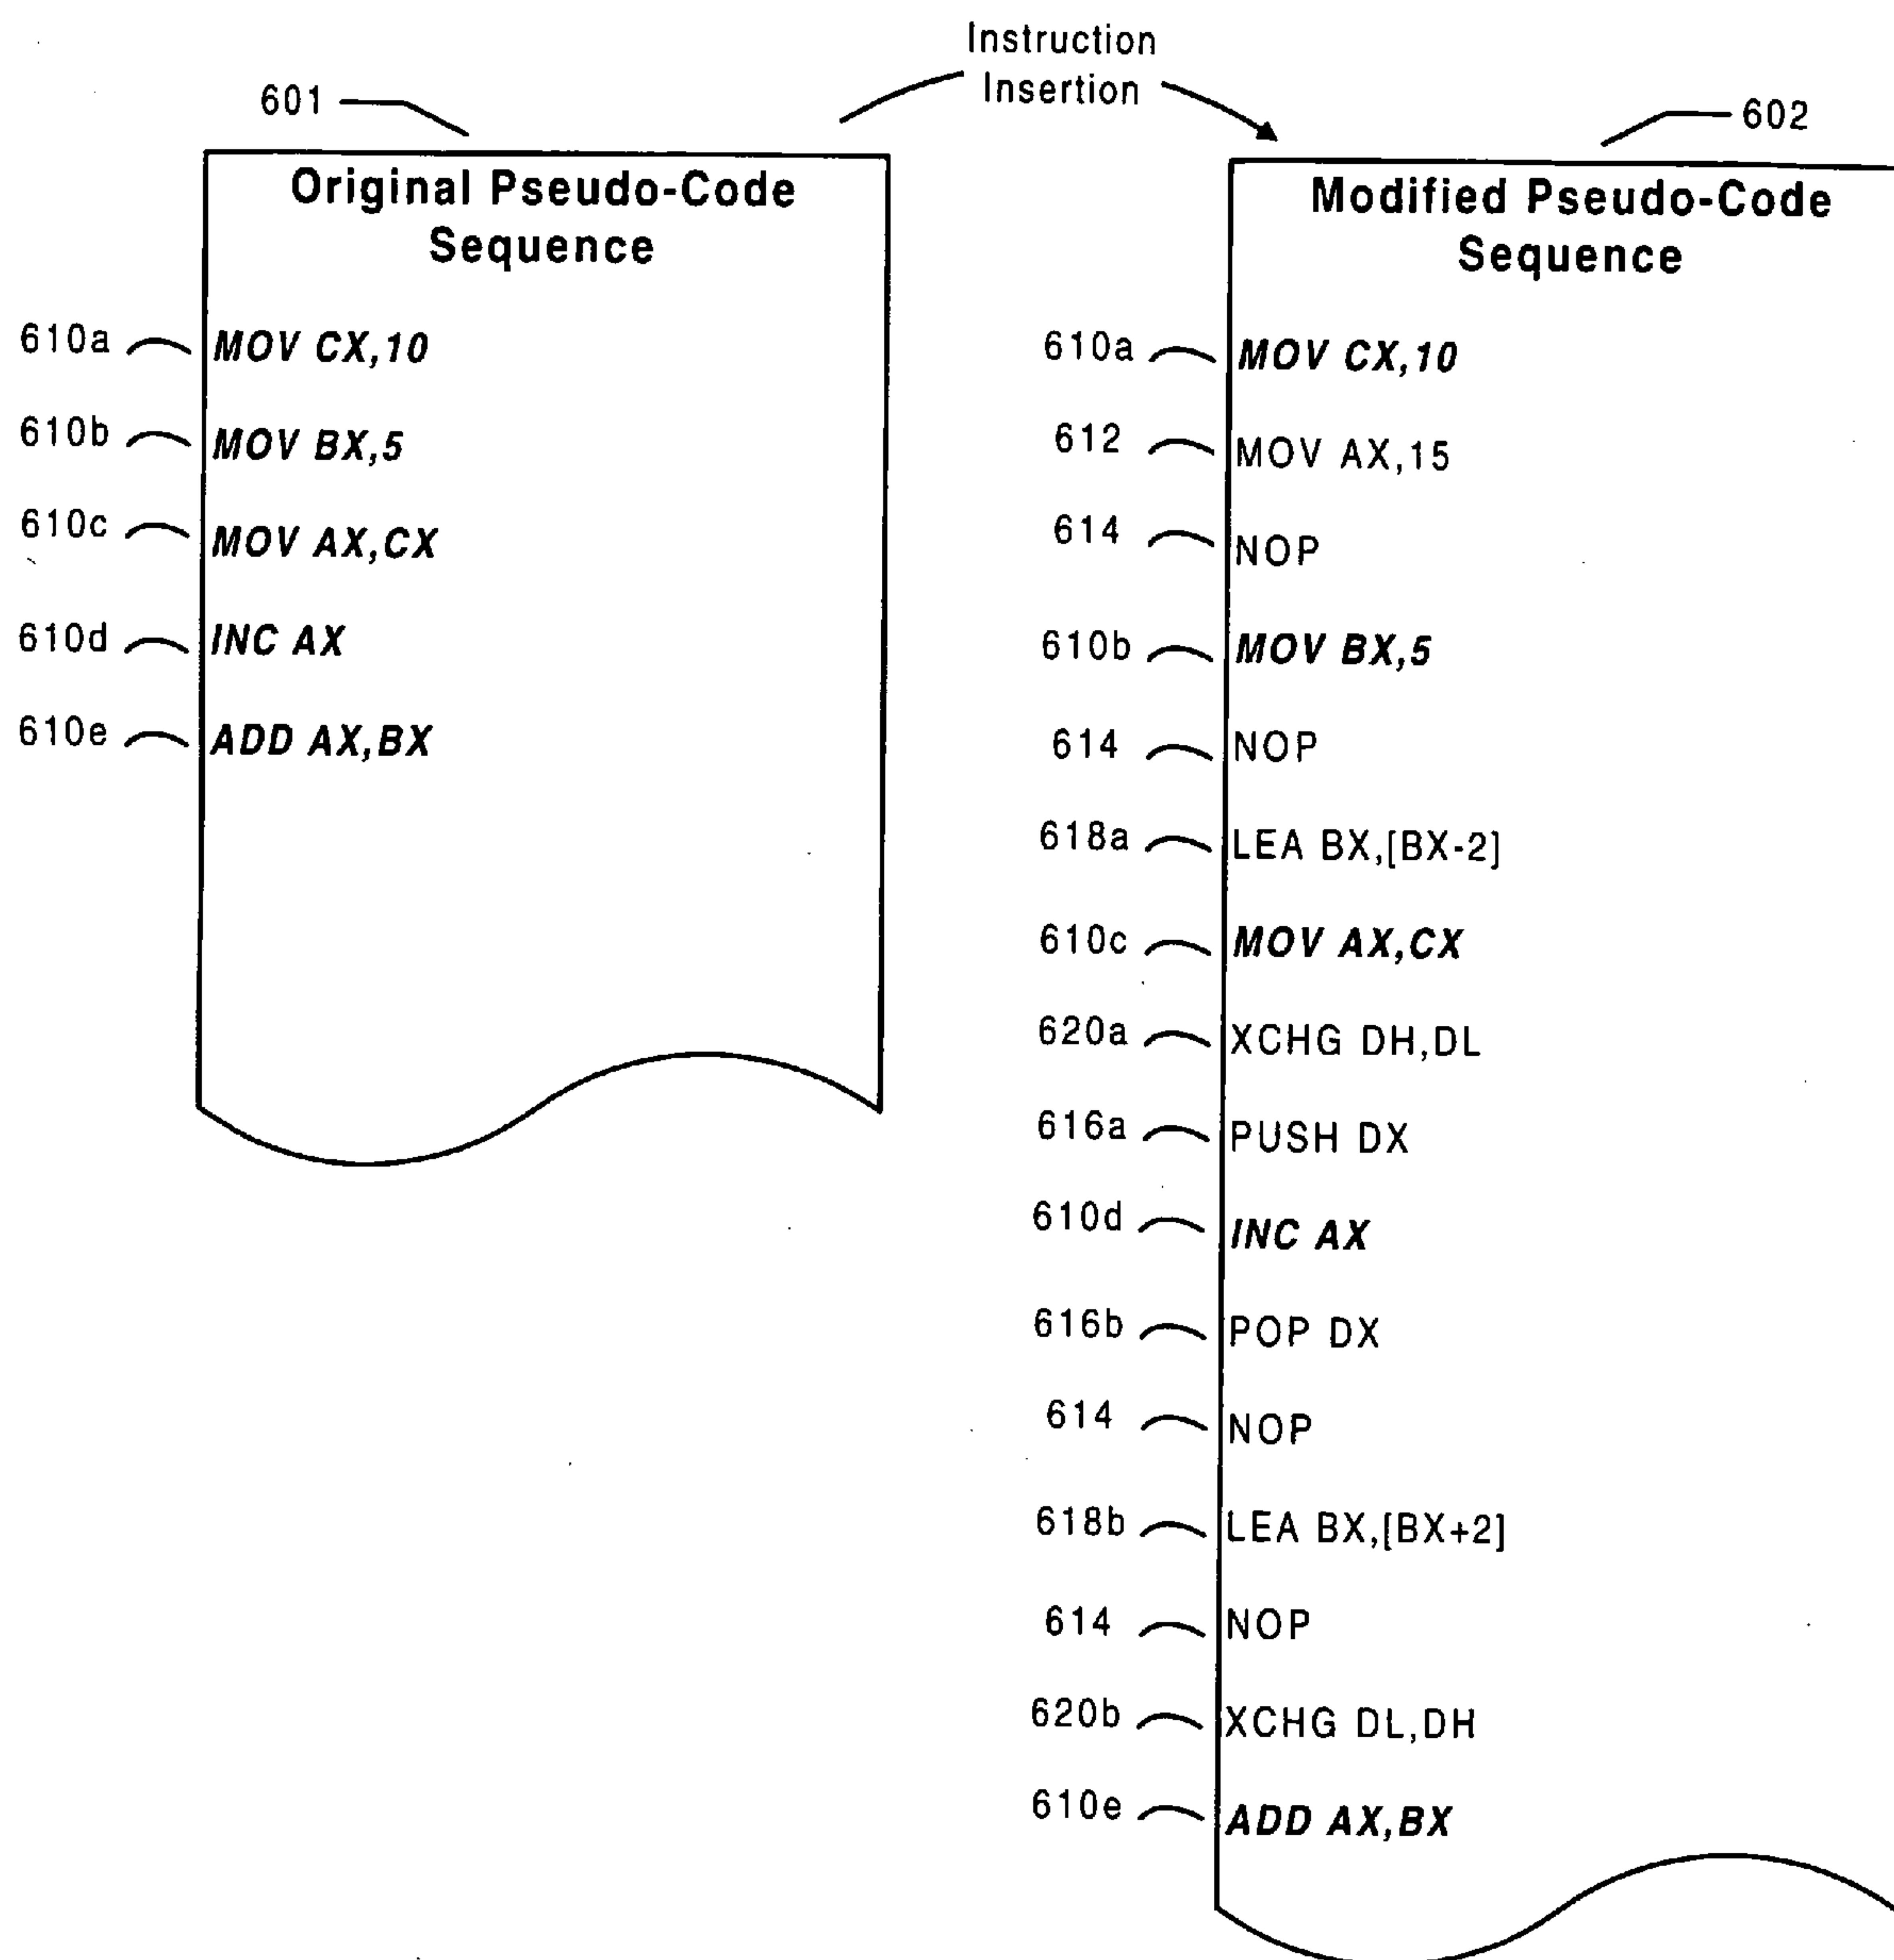


Fig. 6




Fig. 7A

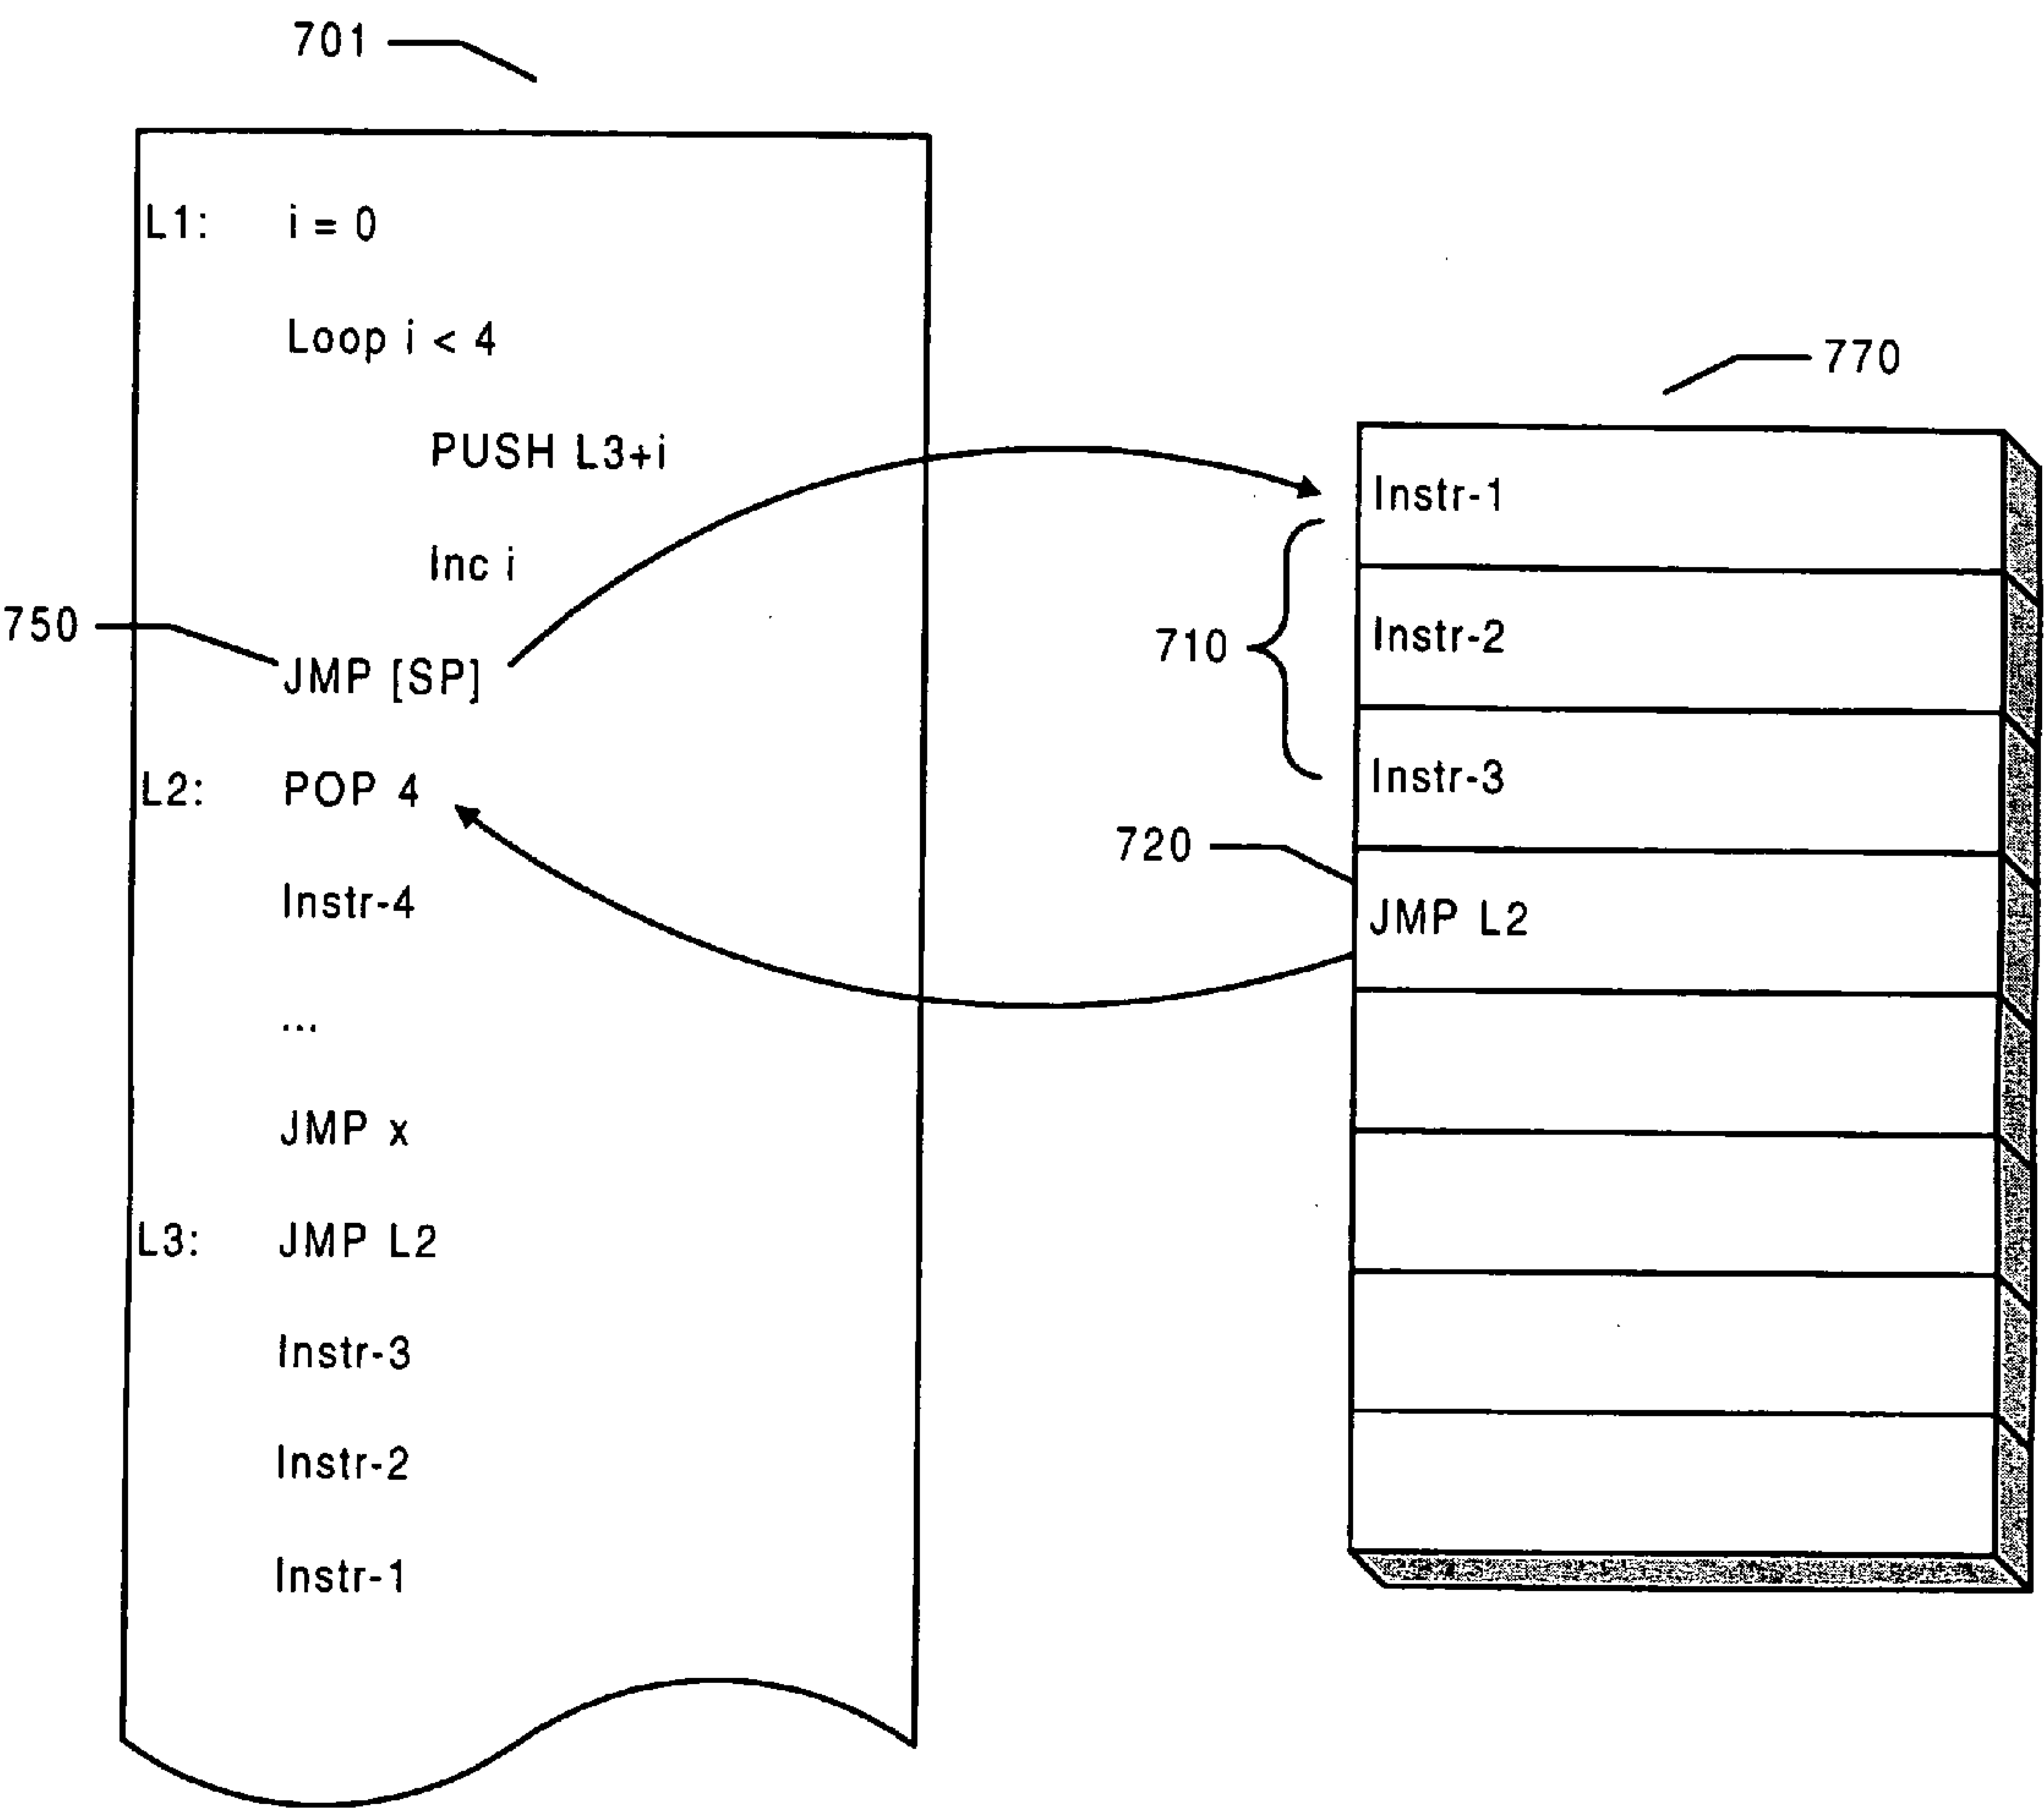


Fig. 7B

800

```

#ifndef CODESPACE_H
#define CODESPACE_H

// A CodeSpace makes it possible to dynamically allocate space in code section.

int const CodeSpaceBlockSize = 4; // Must be greater than sizeof(void*).

#define CodeSpaceUnregistered 0
#define CodeSpaceRegistered 1
#define DataSpace 2

class CodeSpace {
public:
    CodeSpace();

    void *alloc(int bytes);
    // Allocates "bytes" number of chars.
    // (If no code-space is available, calls malloc for data-space.)

    void dealloc(void *block);
    // Deallocates "block" allocated using "alloc".

    inline void create() {
        // Creates space at the point of call and registers it.

        char *space;
        _asm {
            push eax
            lea eax, start
            mov space, eax
            pop eax
            jmp end
            start:
                emit CodeSpaceUnregistered
                // Create space of CodeSpaceBlockSize.
                nop
                nop
                nop
                nop
            end:
        }
        if (*space==CodeSpaceUnregistered) Register(space);
    }

private:
    void *freeBlockList;
    // Linked list of free blocks, inlined in the blocks.

    void Register(char *space);
    // Adds "space" in freeBlockList.
};

#endif /* CODESPACE_H */

```

Fig. 8A

800

```
#include "codespace.h"
#include <stdlib.h>
#include <malloc.h>

CodeSpace::CodeSpace() {
    freeBlockList = 0;
}

void CodeSpace::Register(char *space) {
    *space = CodeSpaceRegistered;
    void *block = space+1;
    // Add block in front of freeBlockList.
    *(void **) block = freeBlockList;
    freeBlockList = block;
}

void *CodeSpace::alloc(int bytes) {
    if (bytes > CodeSpaceBlockSize || freeBlockList == 0) {
        // Fall back to malloc.
        char *space = (char *) malloc(bytes+1);
        if (space==0) return 0; // malloc failed
        *space = DataSpace;
        return (void *) (space+1);
    }
    // Return the block from the front of the list.
    void *block = freeBlockList;
    freeBlockList = *(void **)block;
    return block;
}

void CodeSpace::dealloc(void *block) {
    char *space = ((char *)block)-1;
    if (*space == DataSpace) {
        free(space);
        return;
    }
    *(void **) block = freeBlockList;
    freeBlockList = block;
}
```

Fig. 8B

```
#include <stdio.h>
#include "codespace.h"

CodeSpace cs;

void main() {

    // Create some blocks
    830  cs.create();
        cs.create();

        int size = CodeSpaceBlockSize;
    832  void *b1 = cs.alloc(size);           // allocate code-space
    834  void *b2 = cs.alloc(size+1);        // allocate using malloc
    836  void *b3 = cs.alloc(size-1);        // allocate code-space
    838  void *b4 = cs.alloc(size);          // allocate using malloc

    840  cs.dealloc(b1);

    842  void *b6 = cs.alloc(size);           // allocate code-space
}
```

Fig. 8C




Fig. 8D

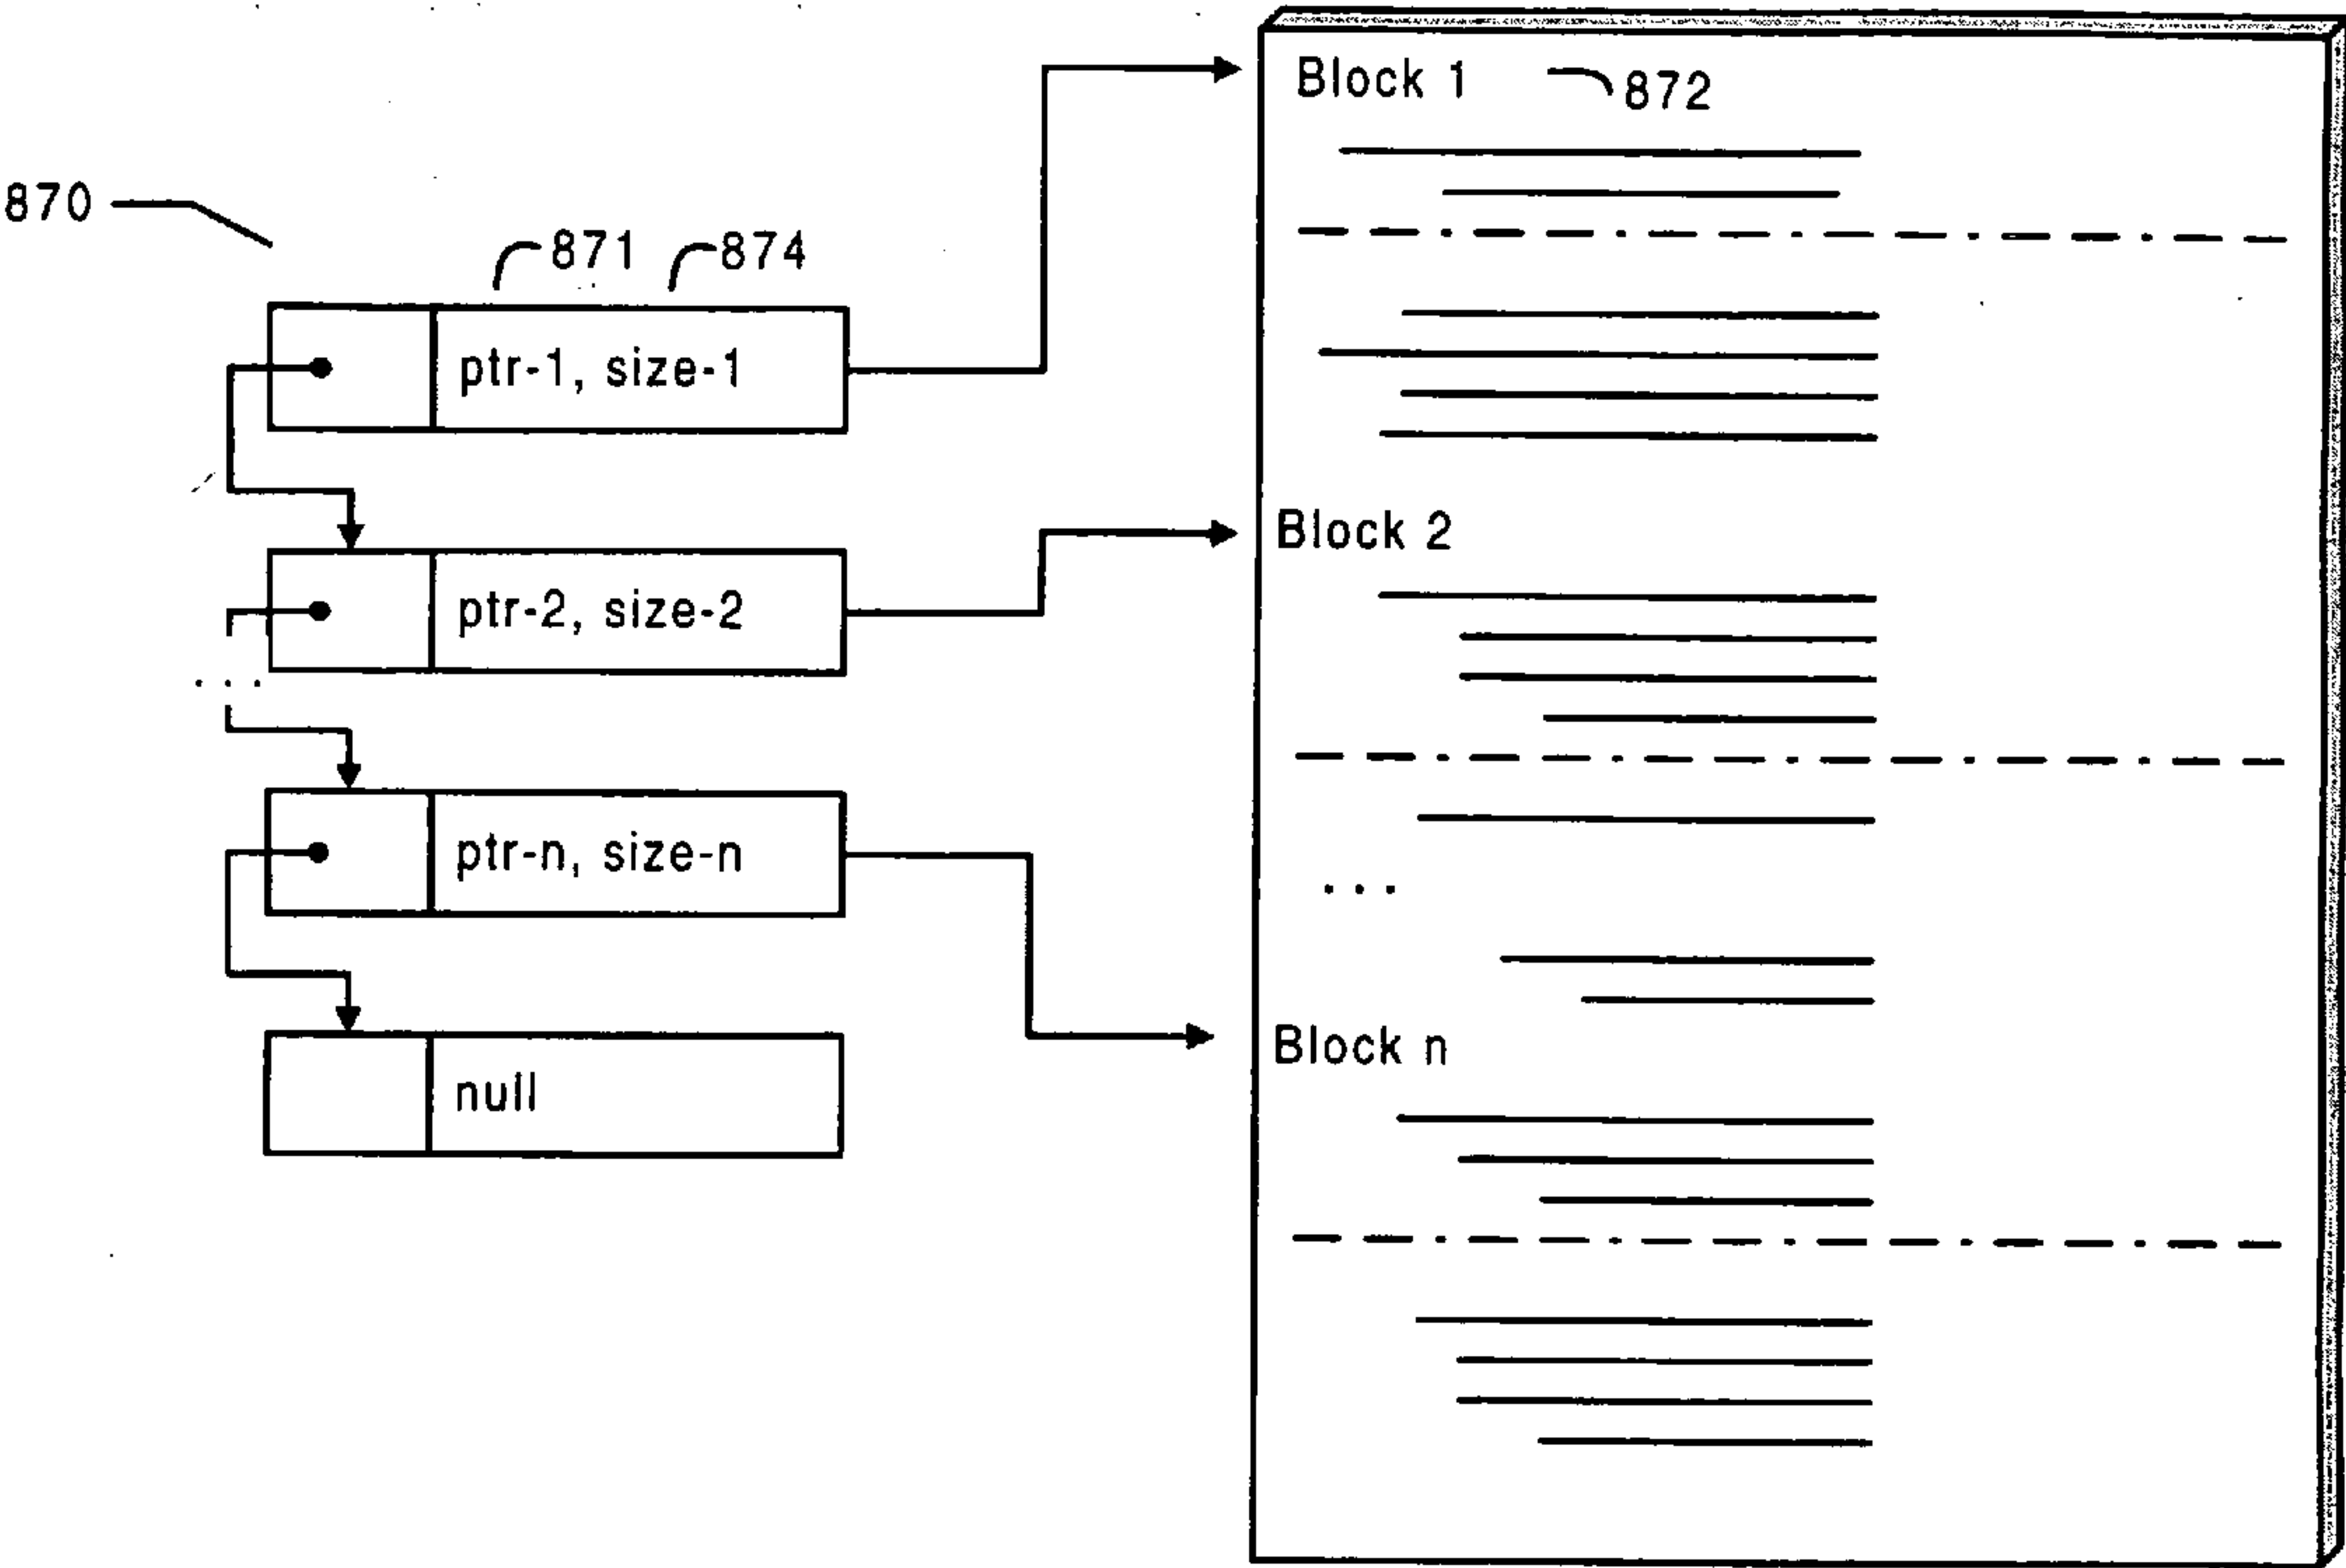


Fig. 8E

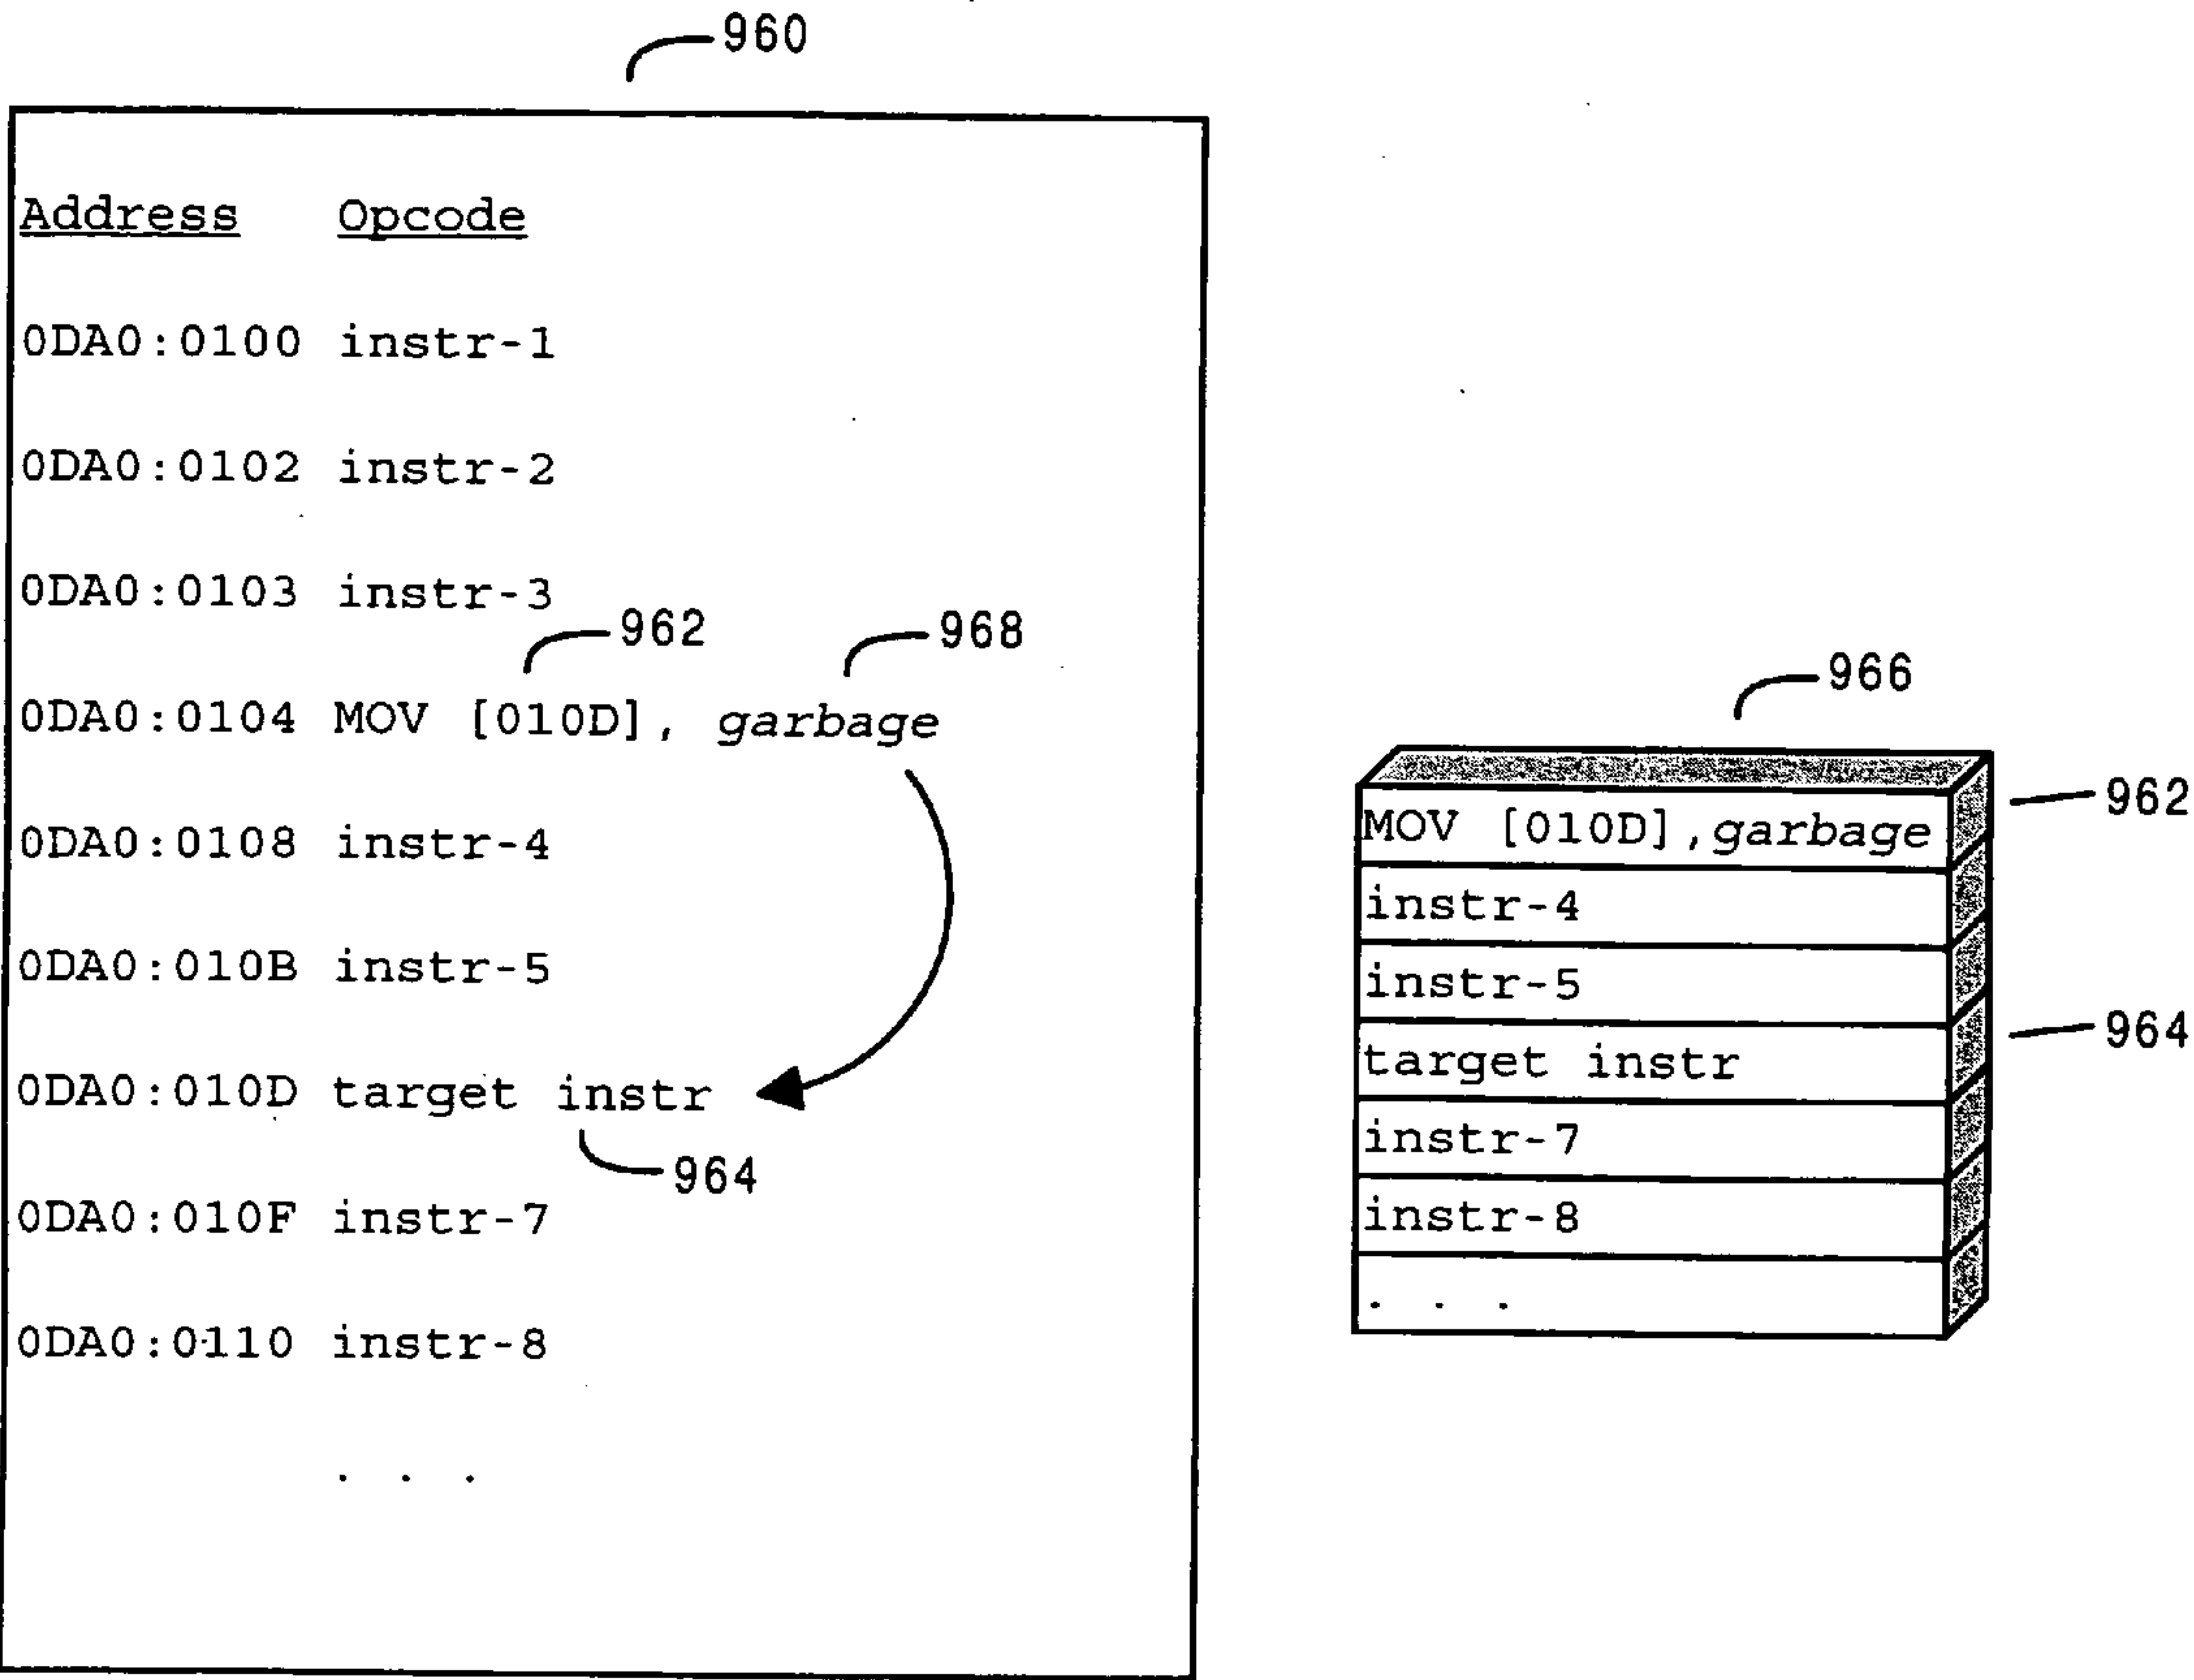


Fig. 9

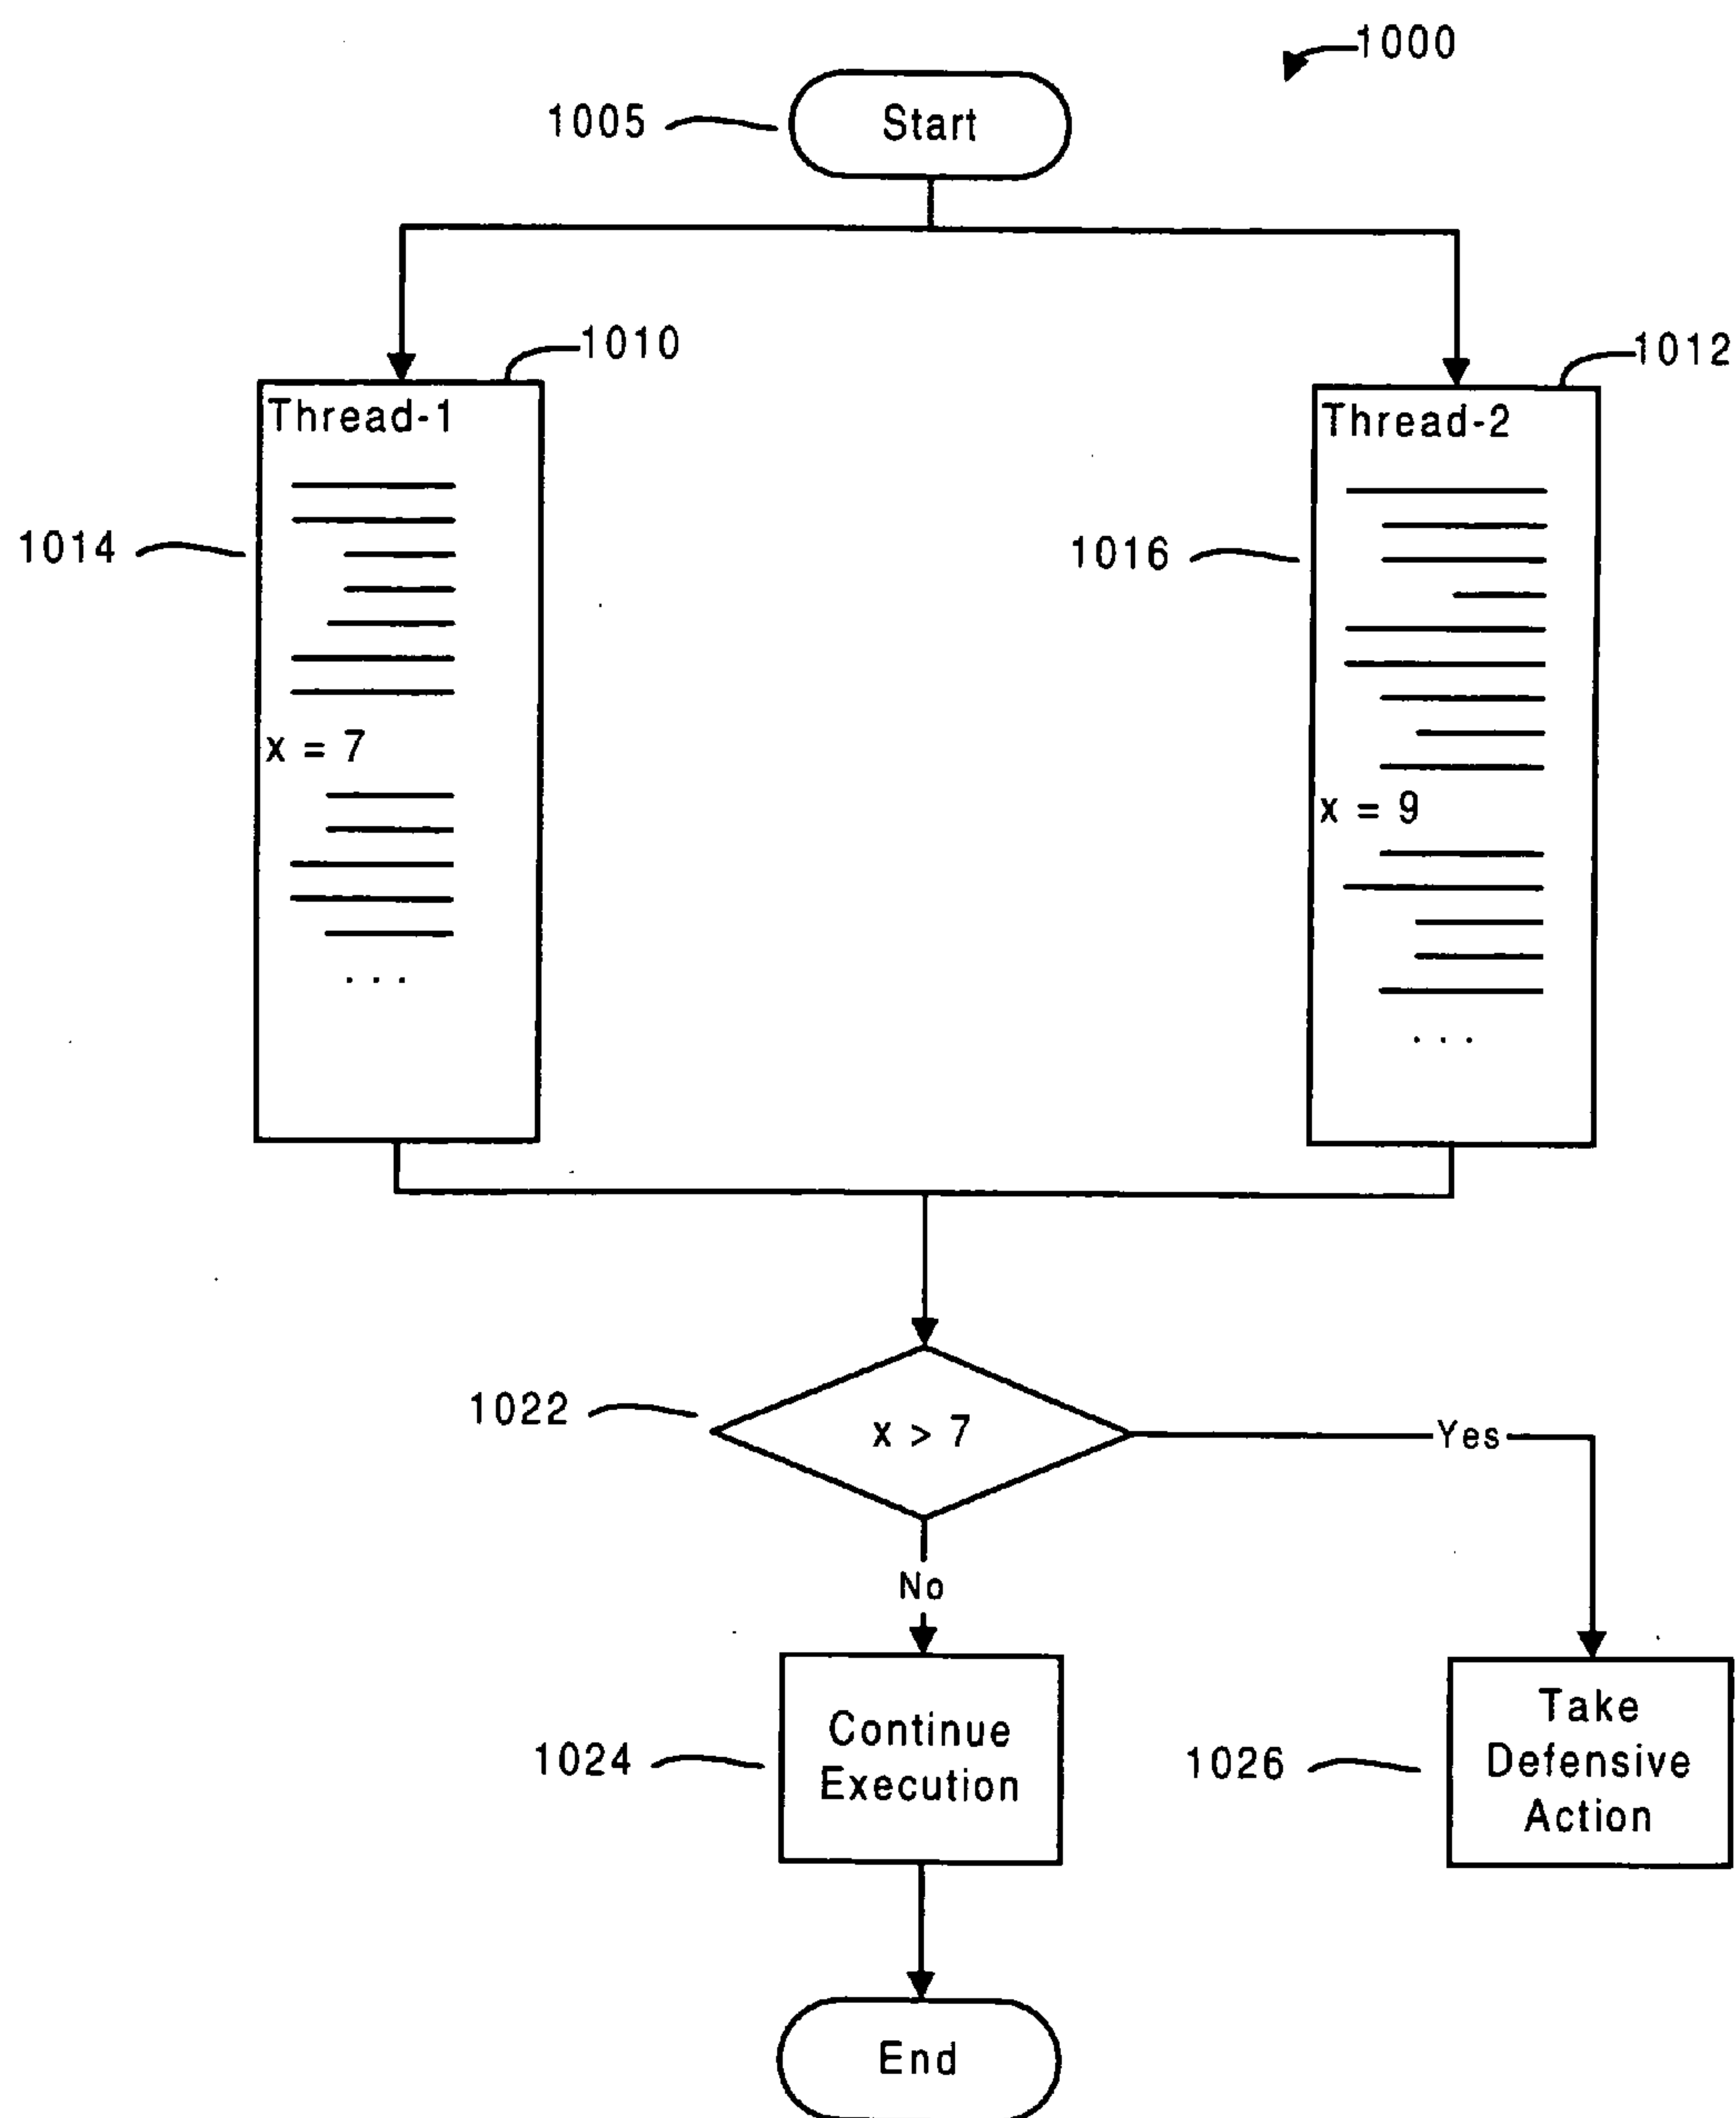


FIG. 10

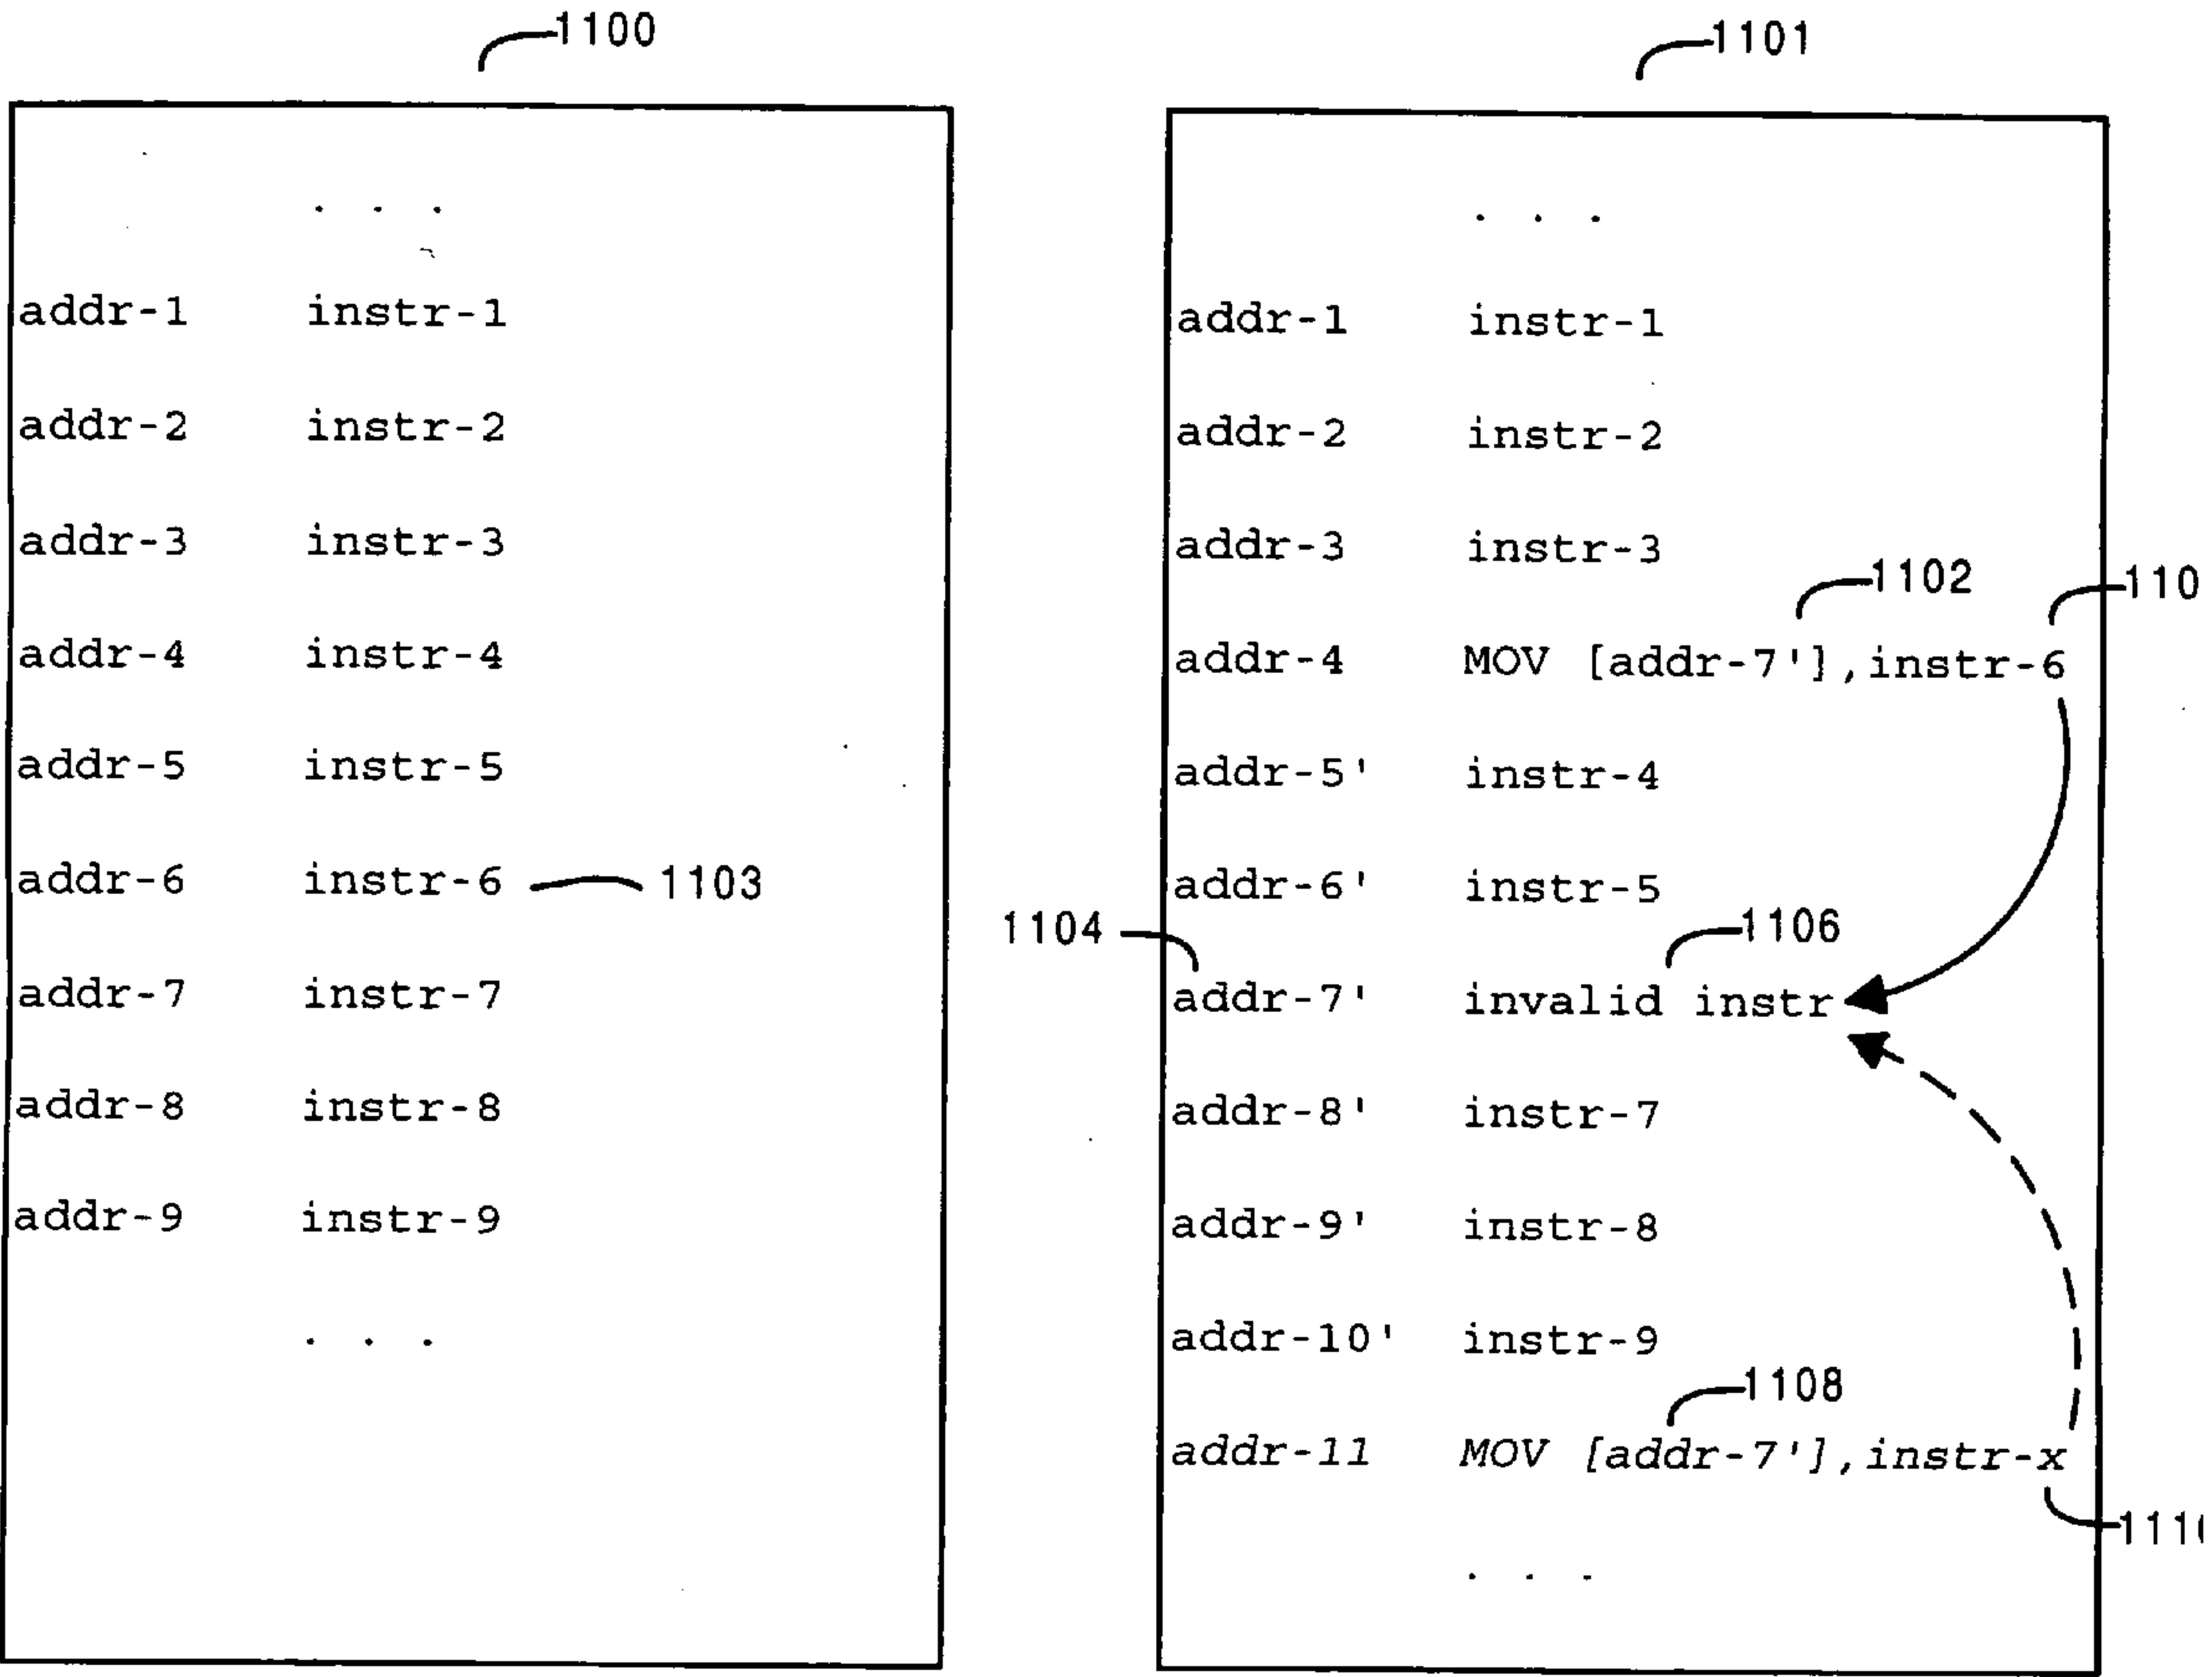


Fig. 11

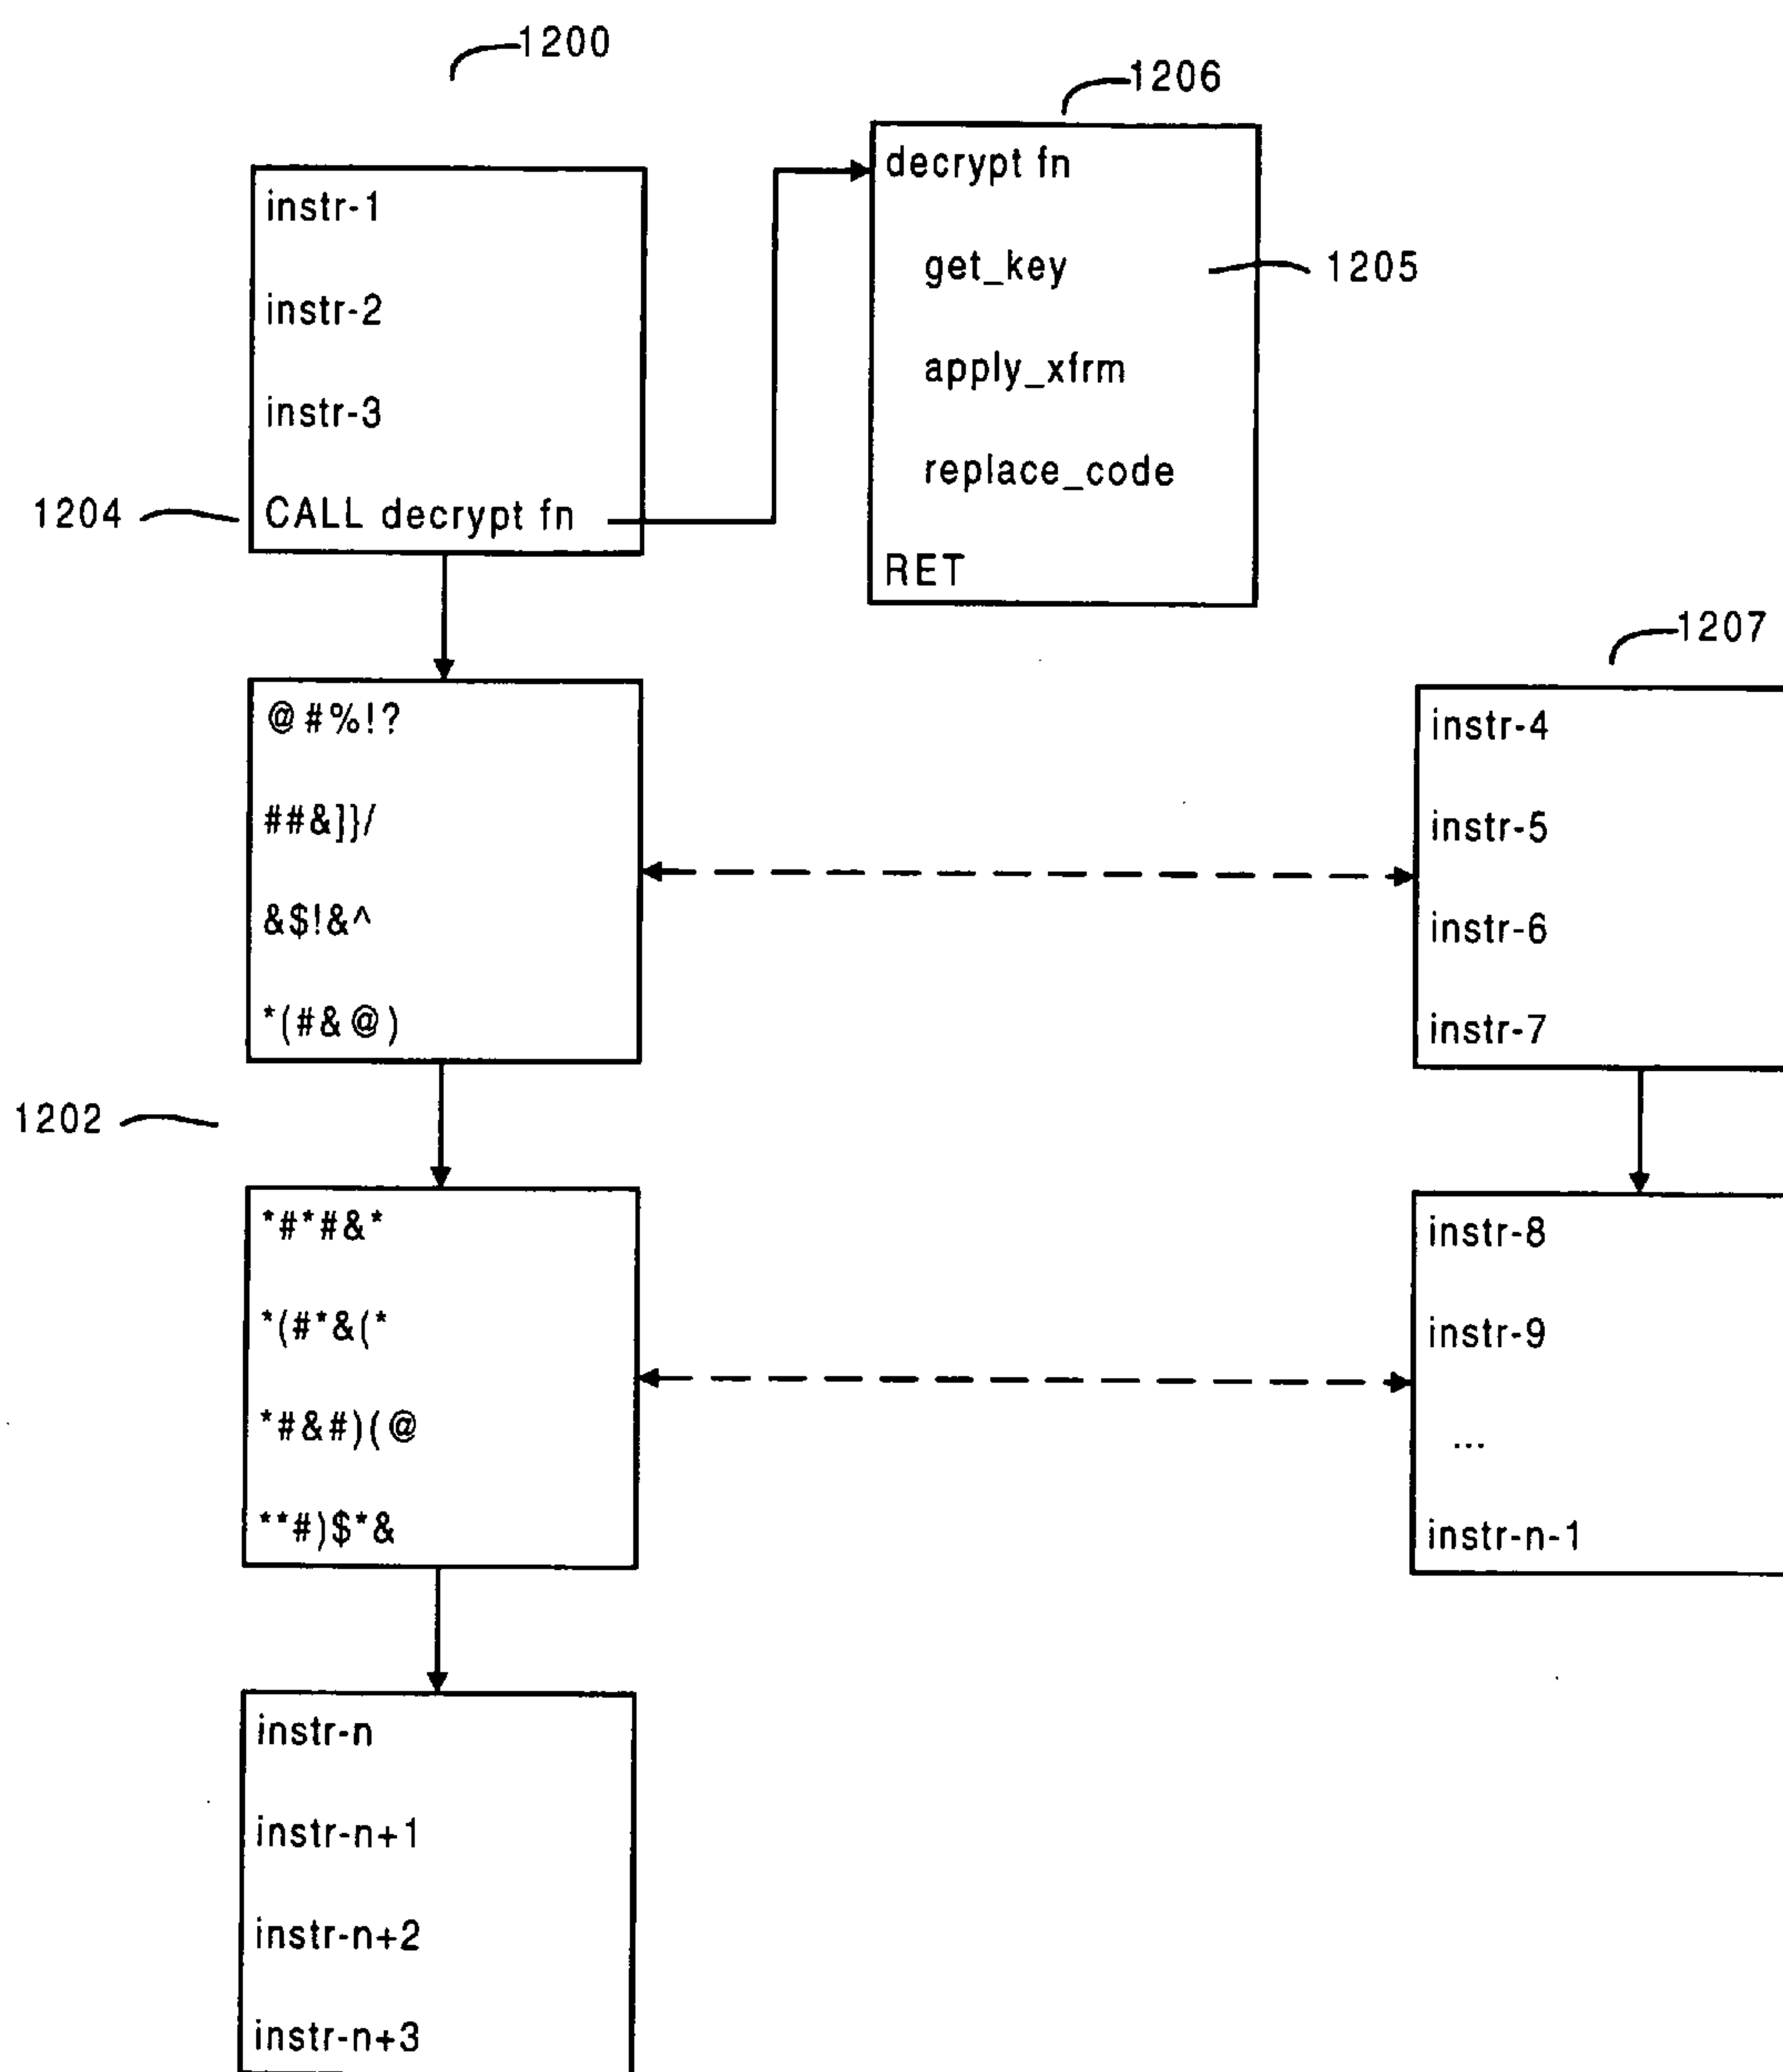


Fig. 12

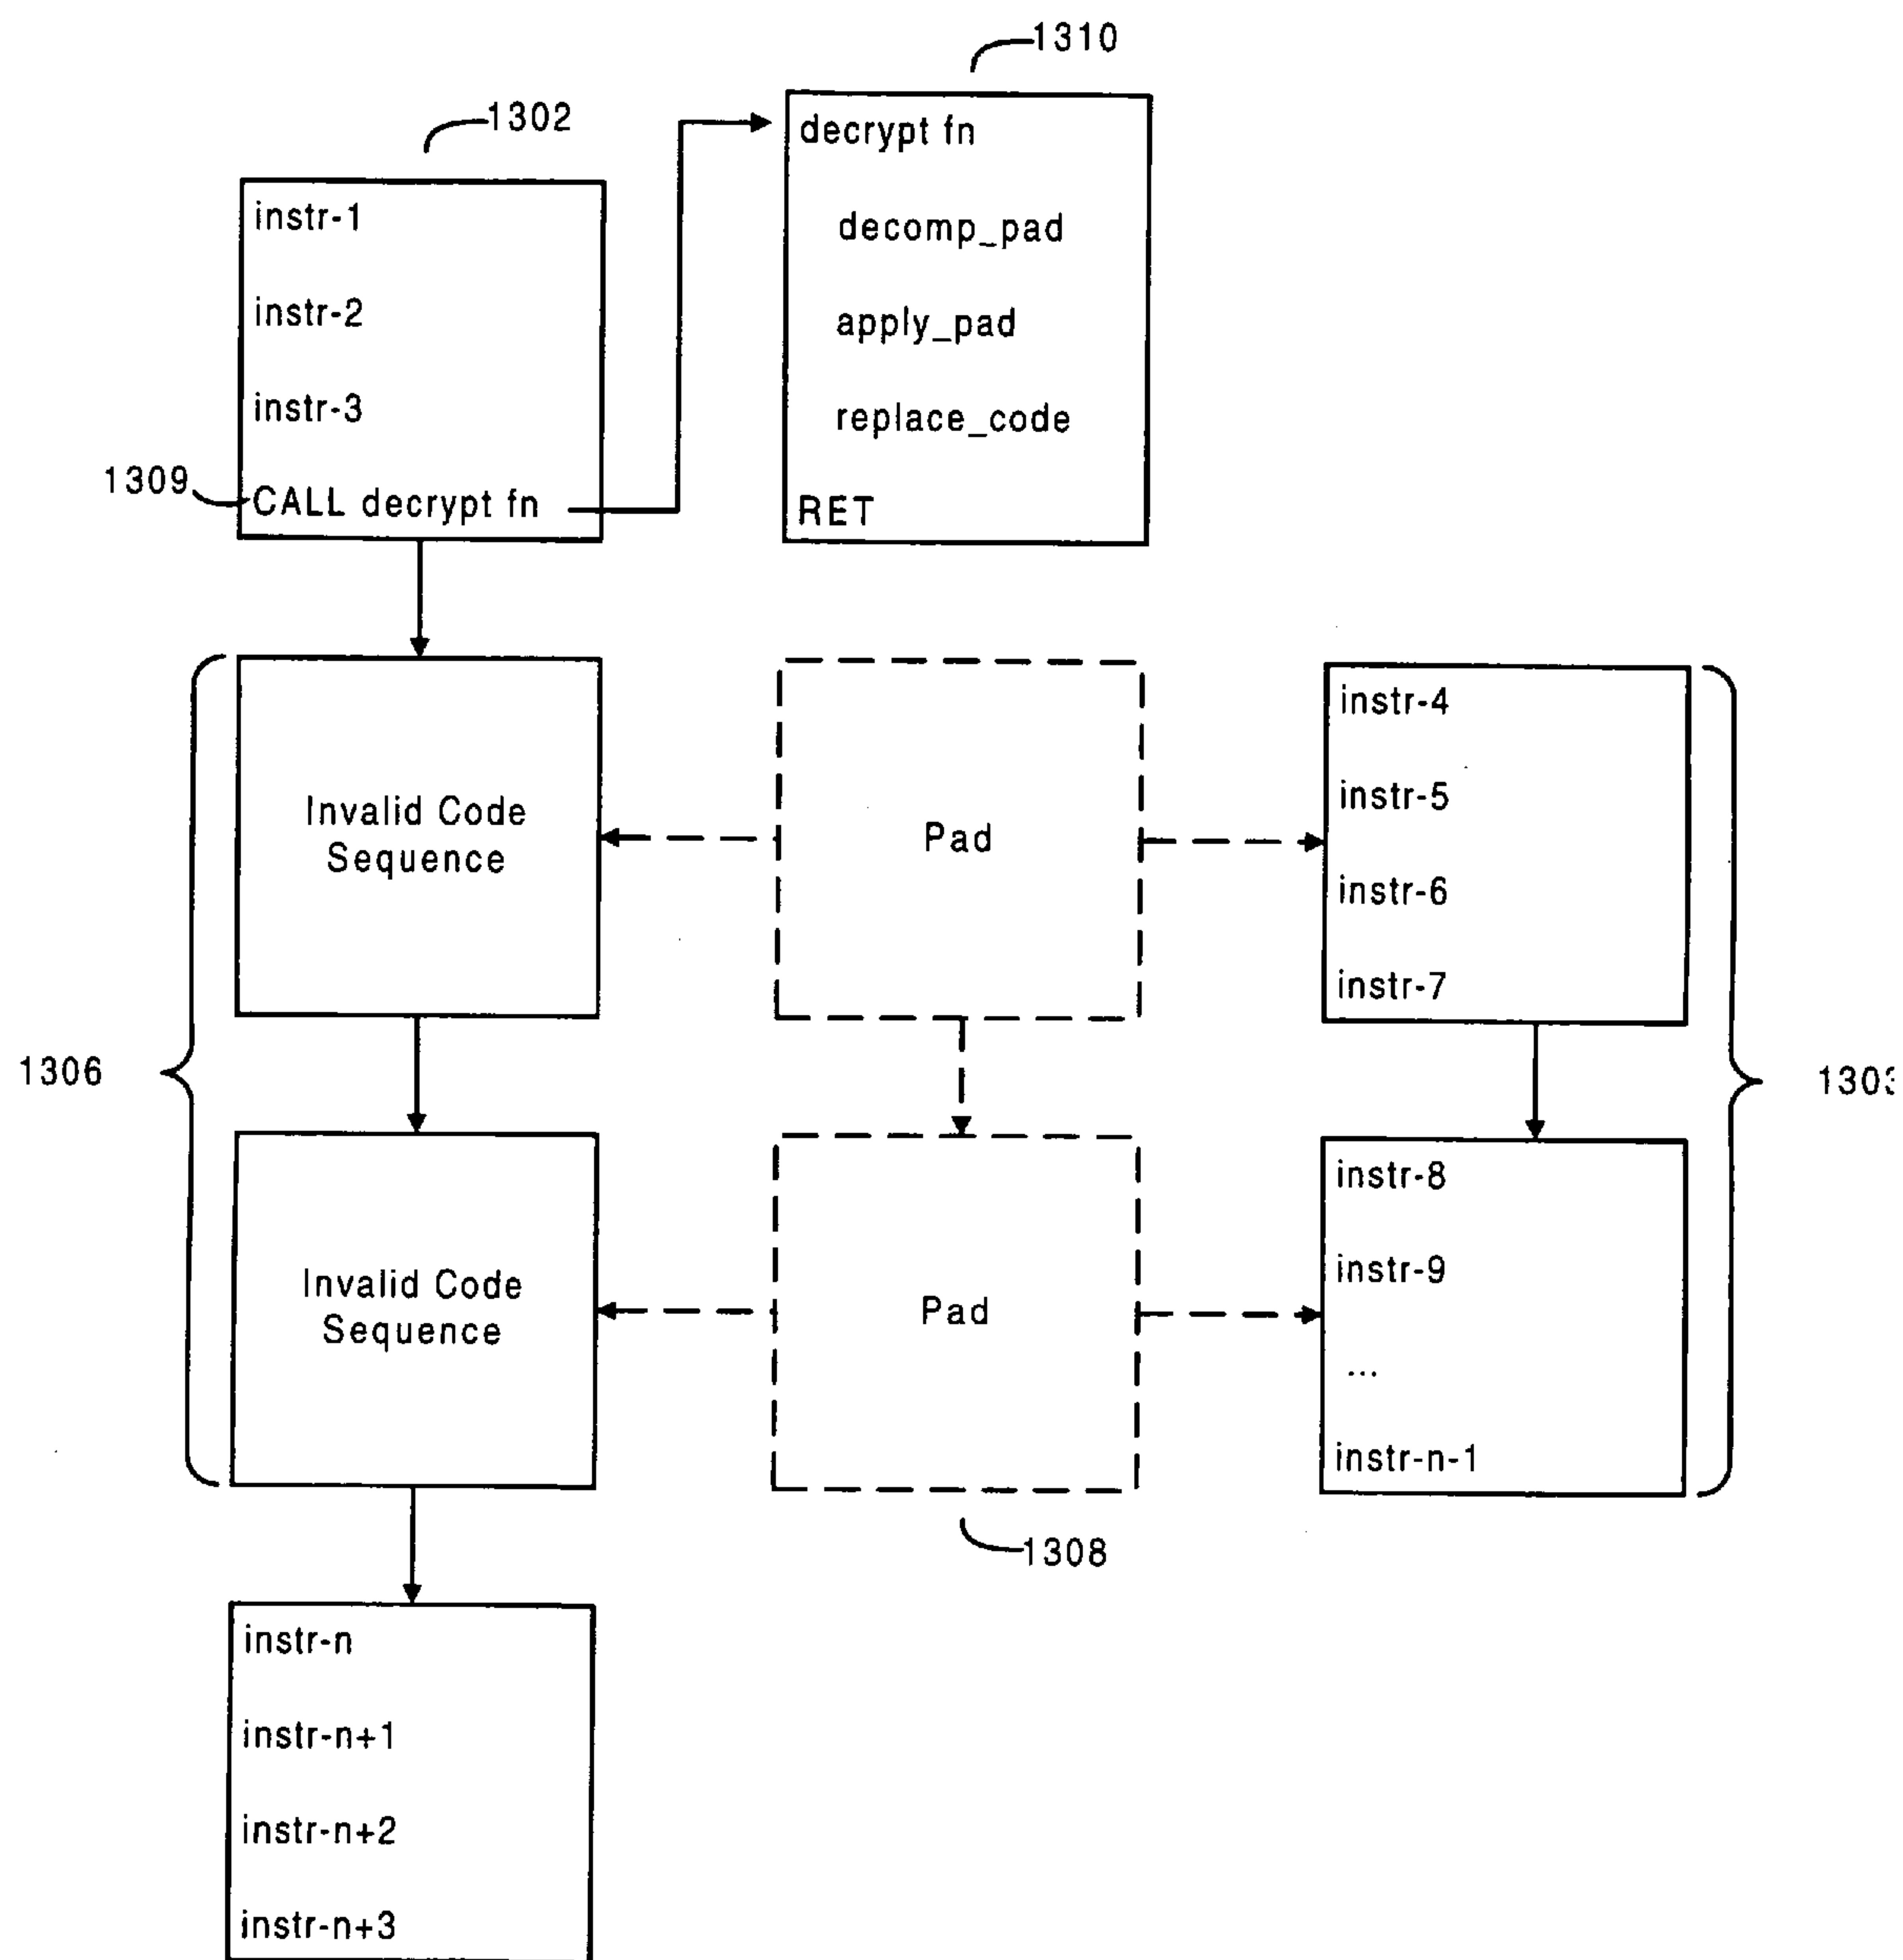


Fig. 13

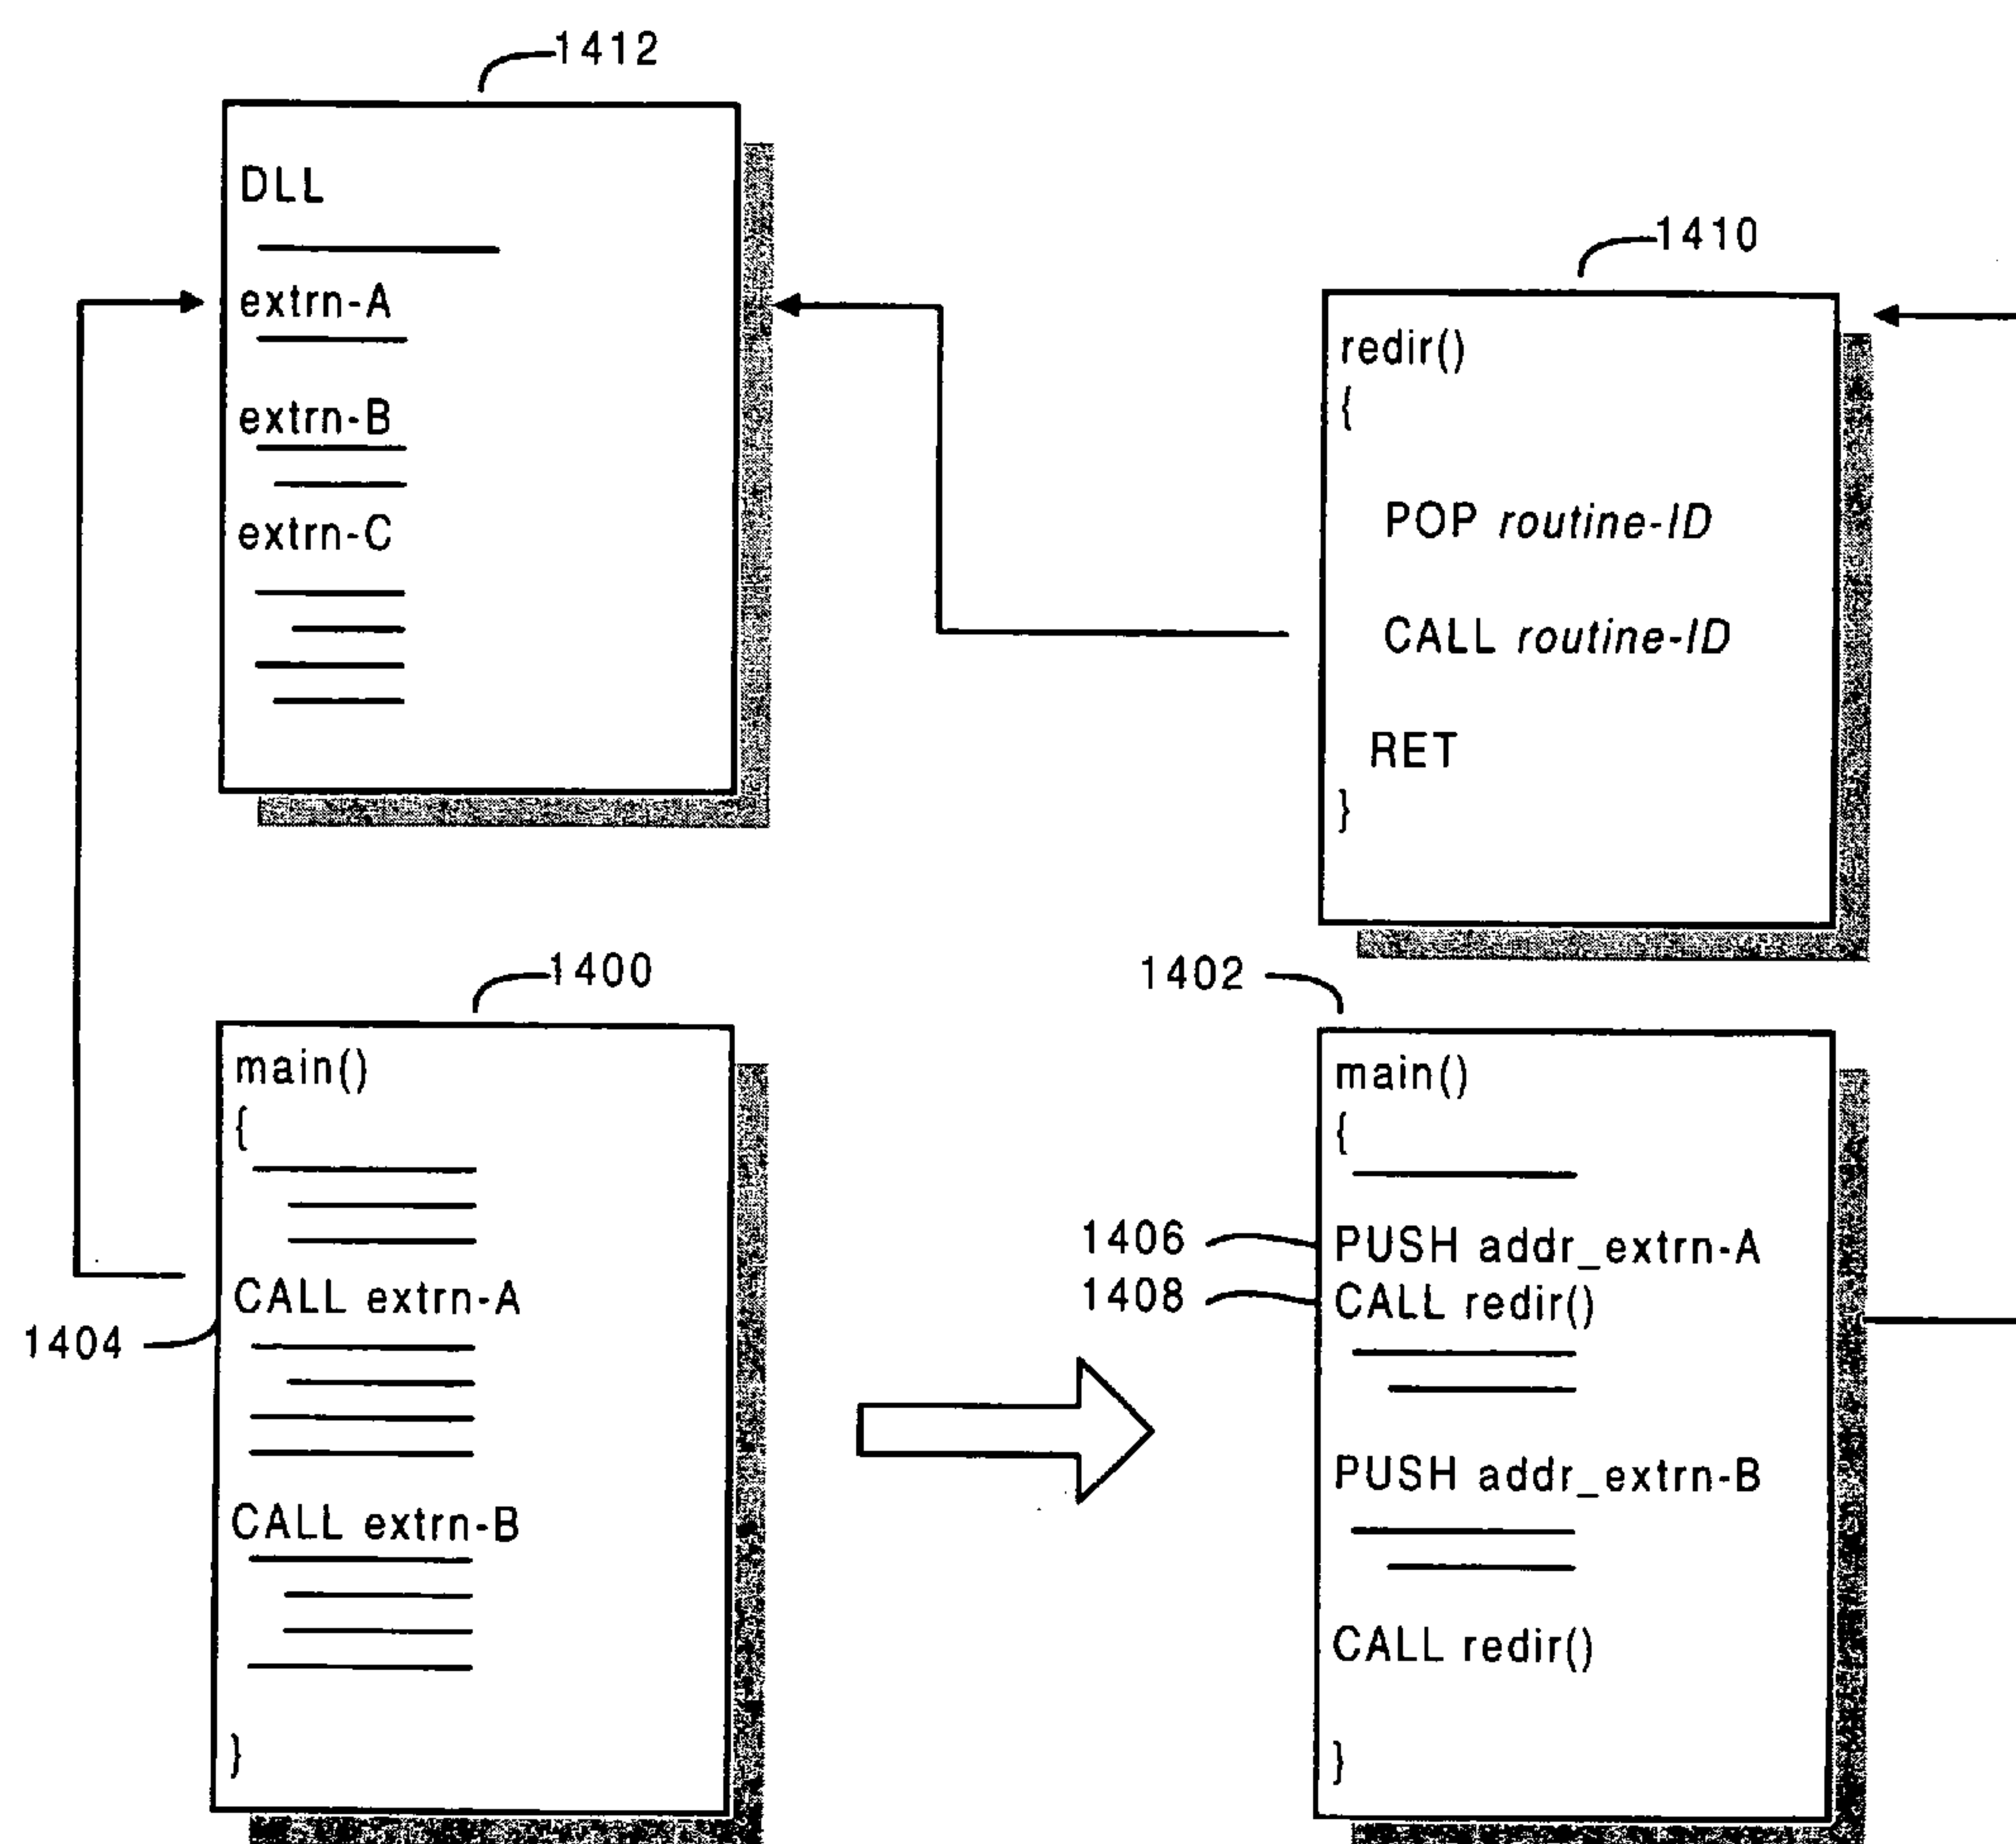


Fig. 14

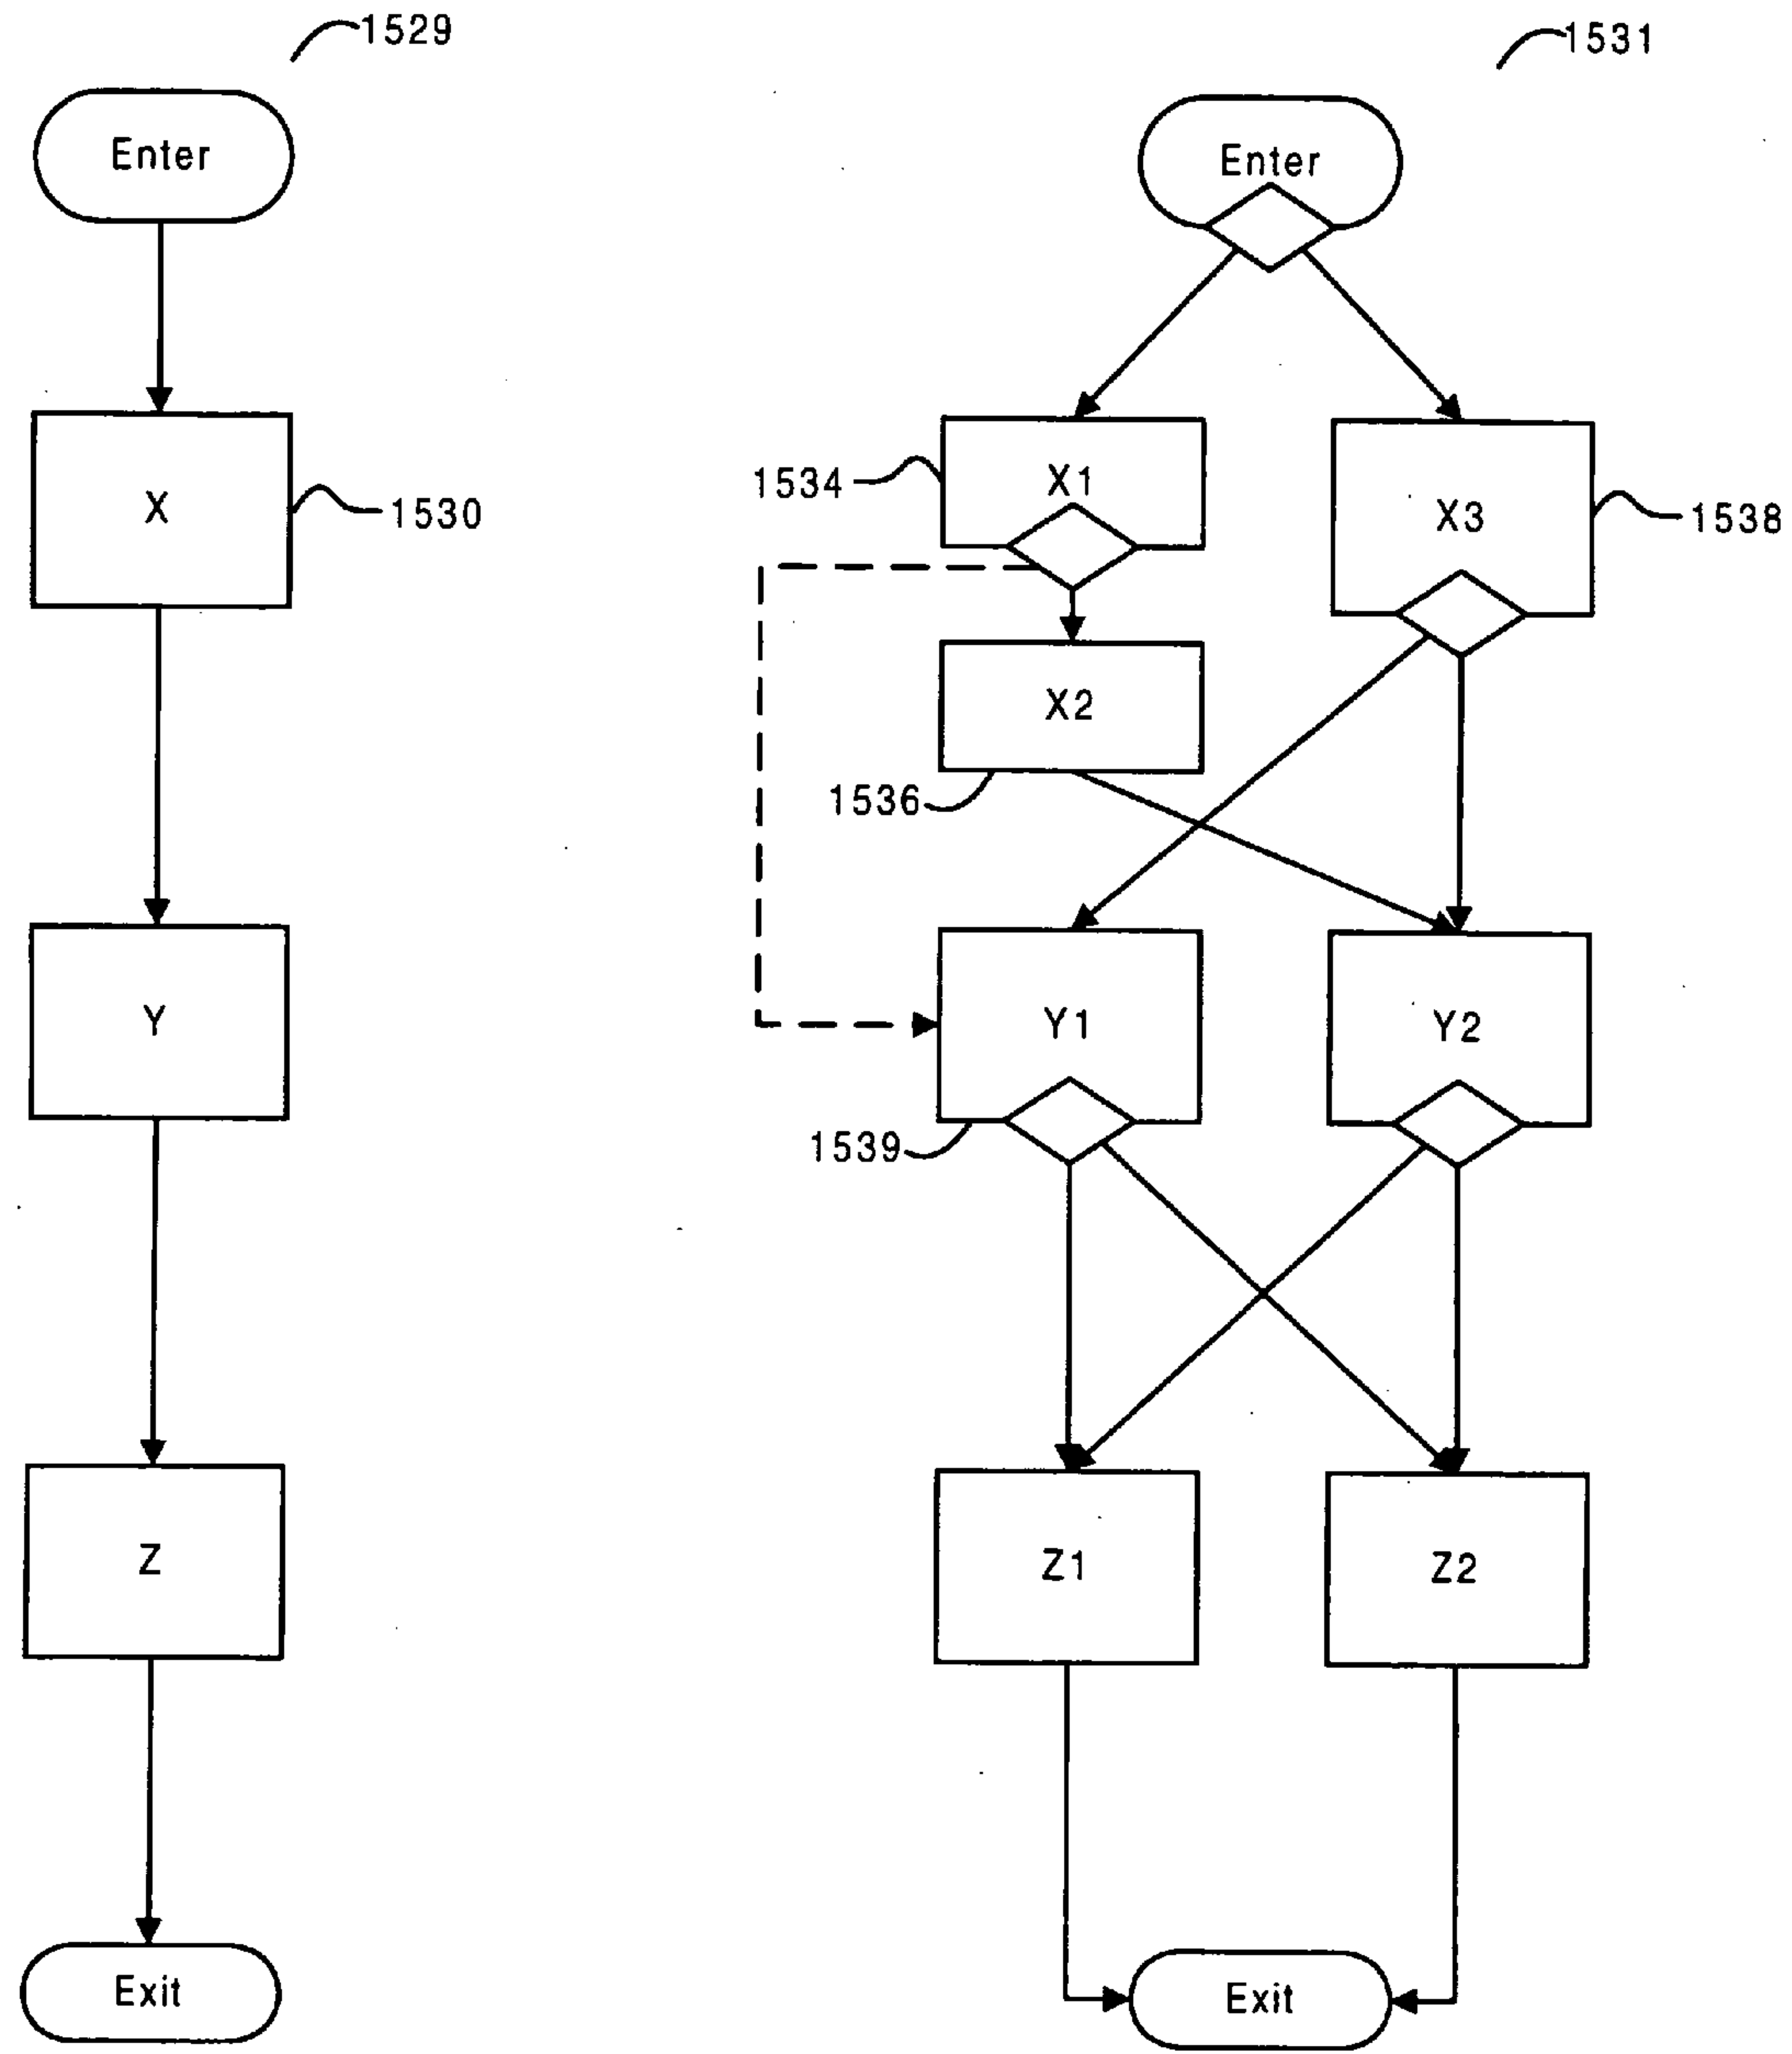


Fig. 15

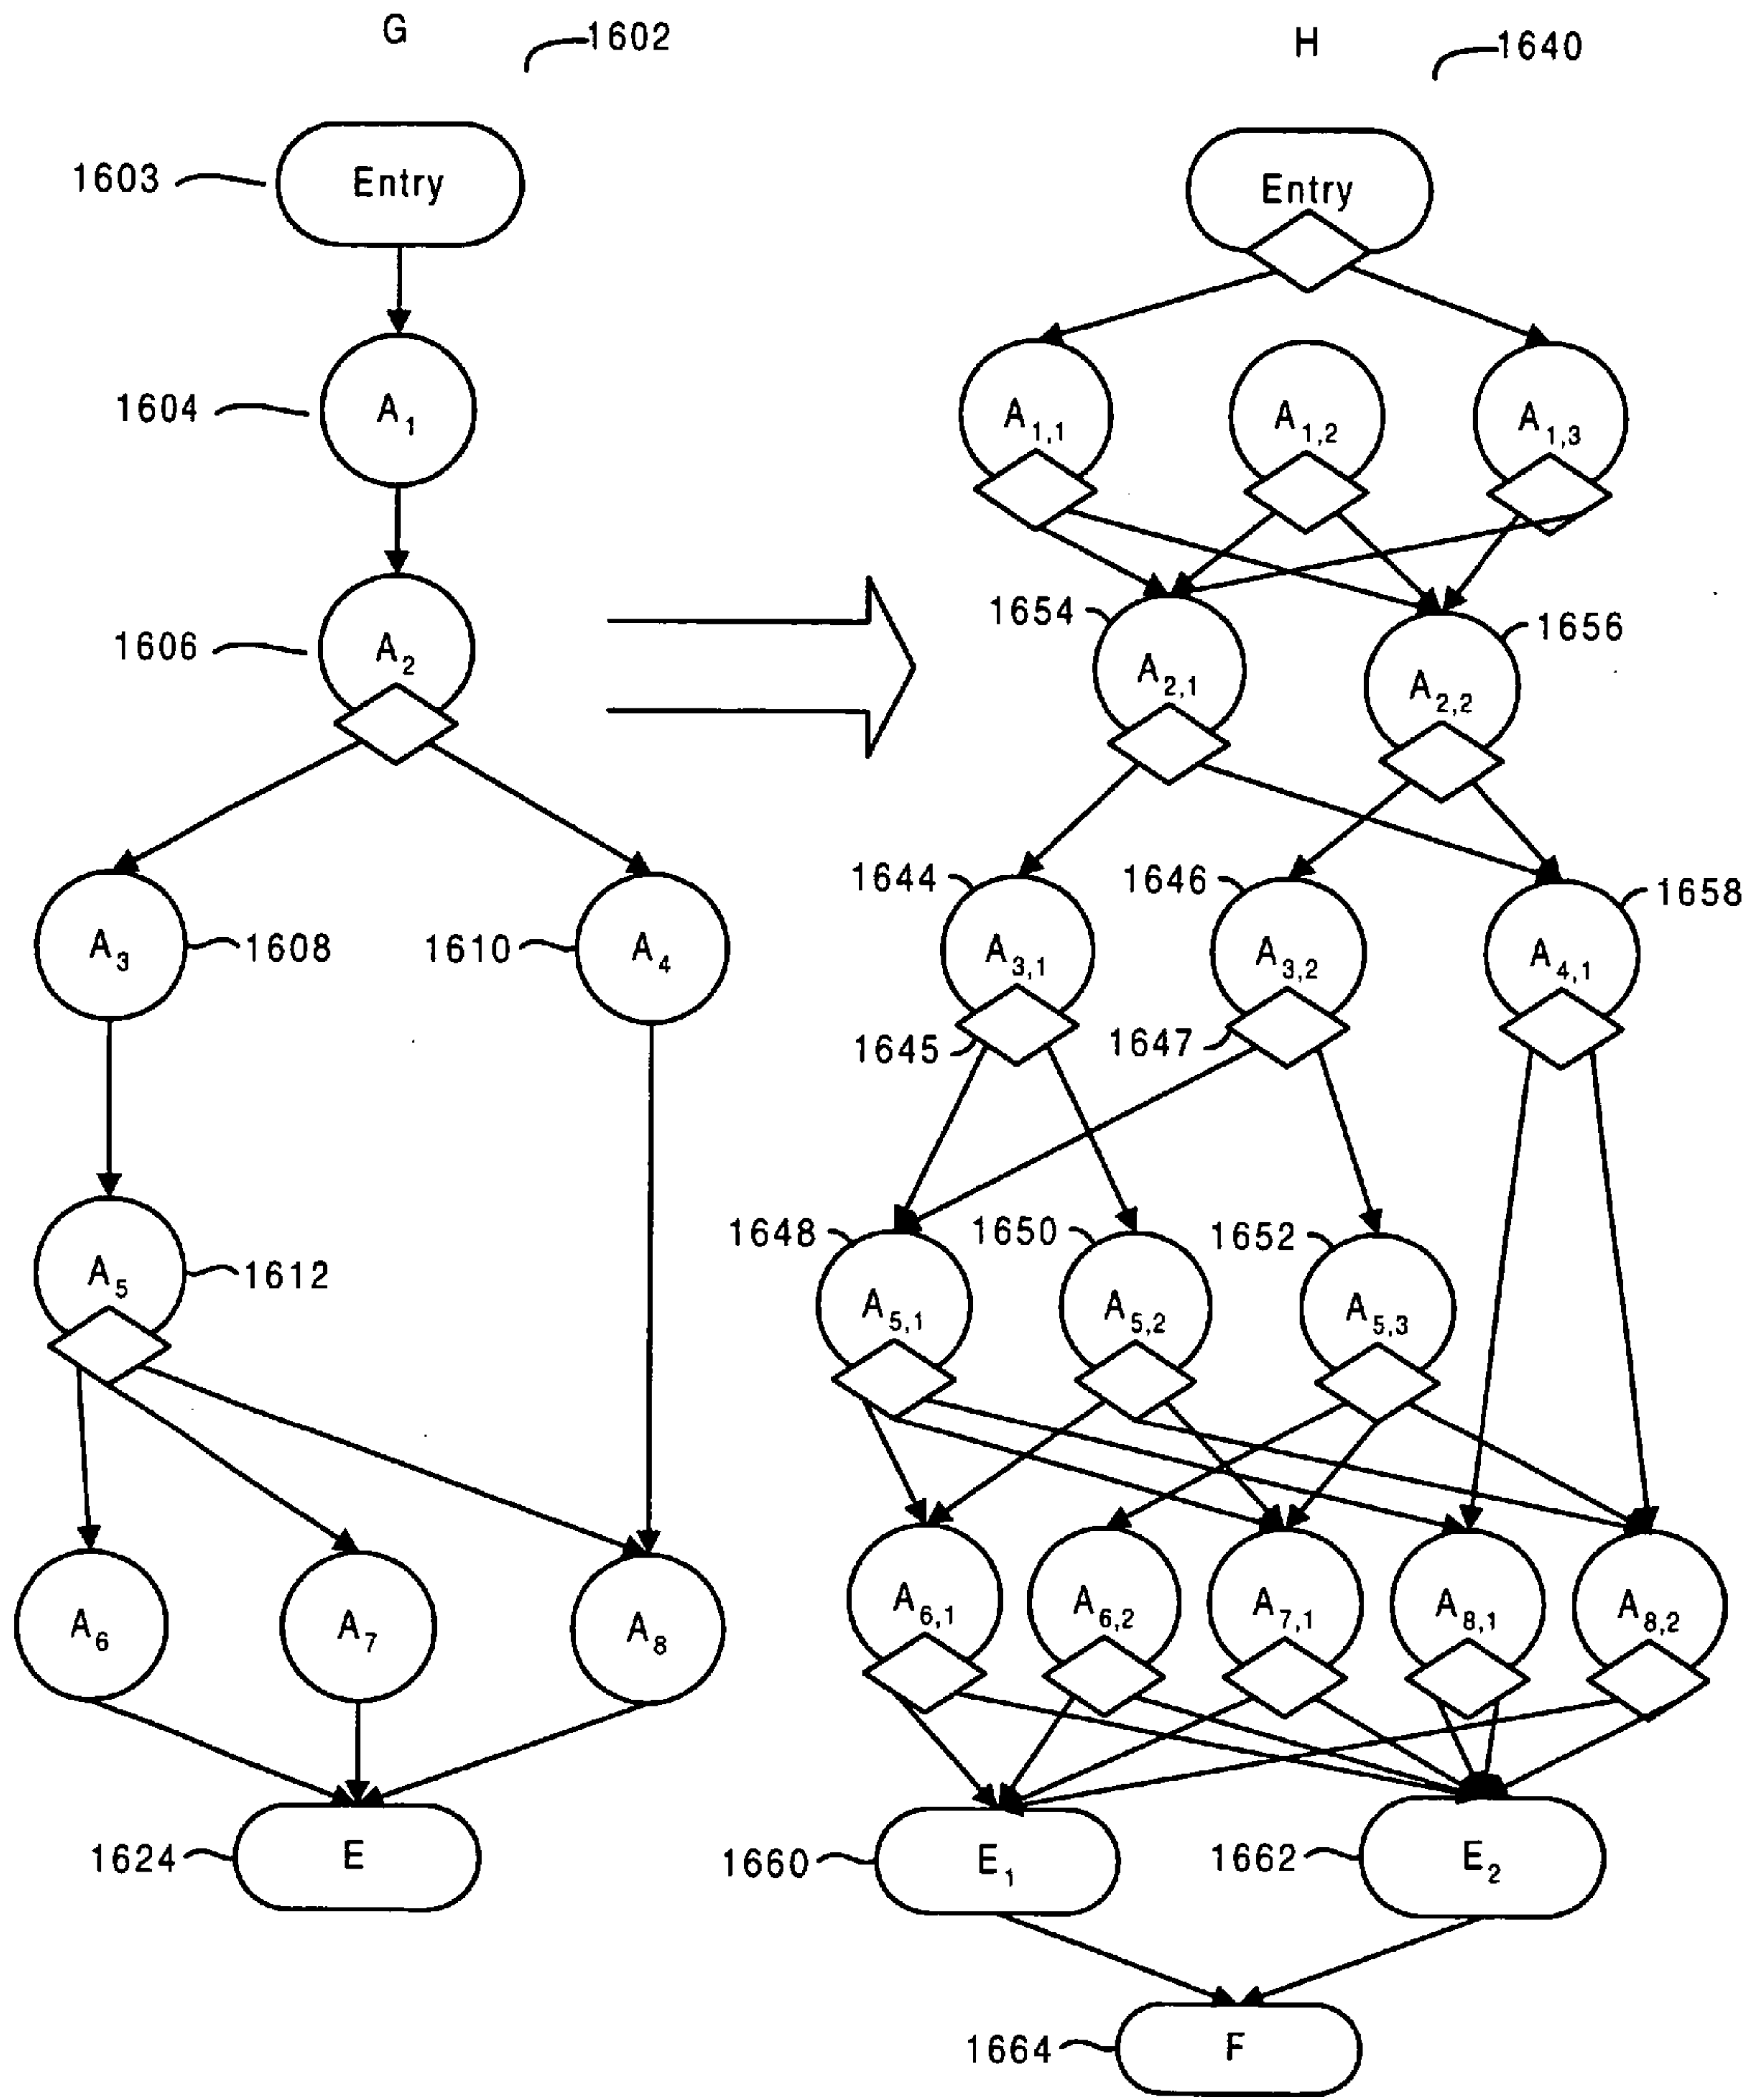


FIG. 16A

FIG. 16B

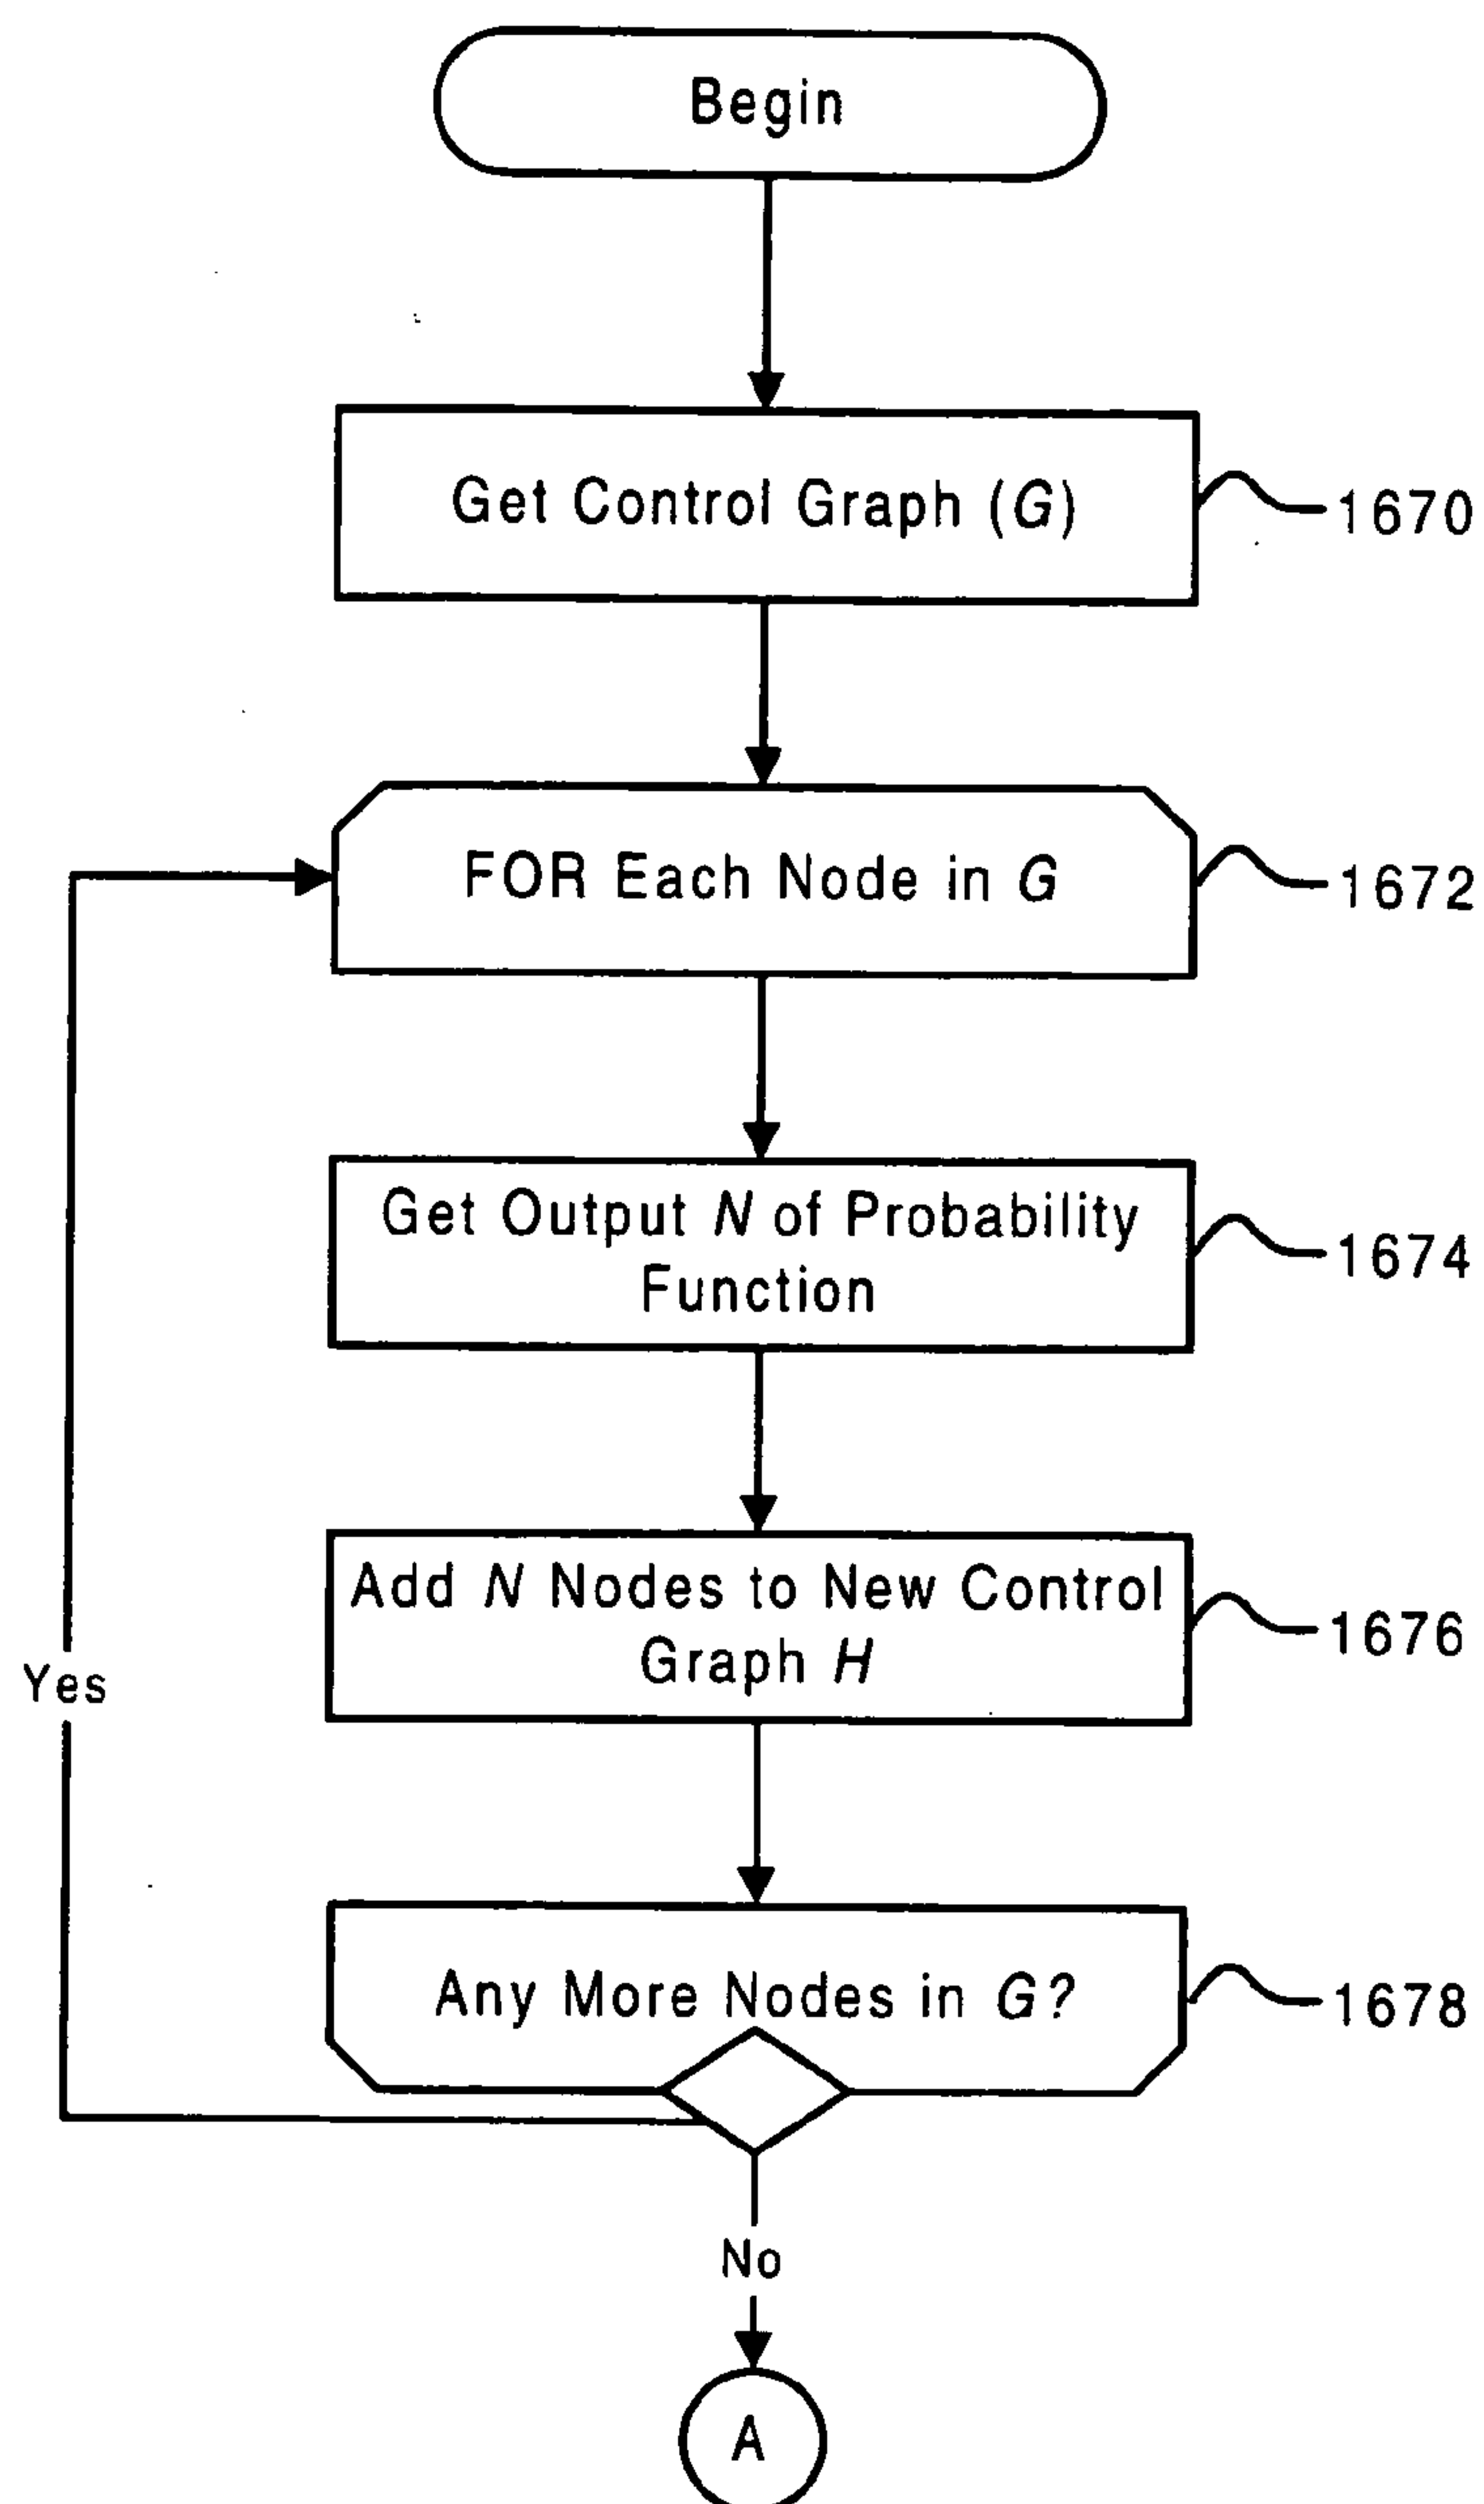


FIG. 16C

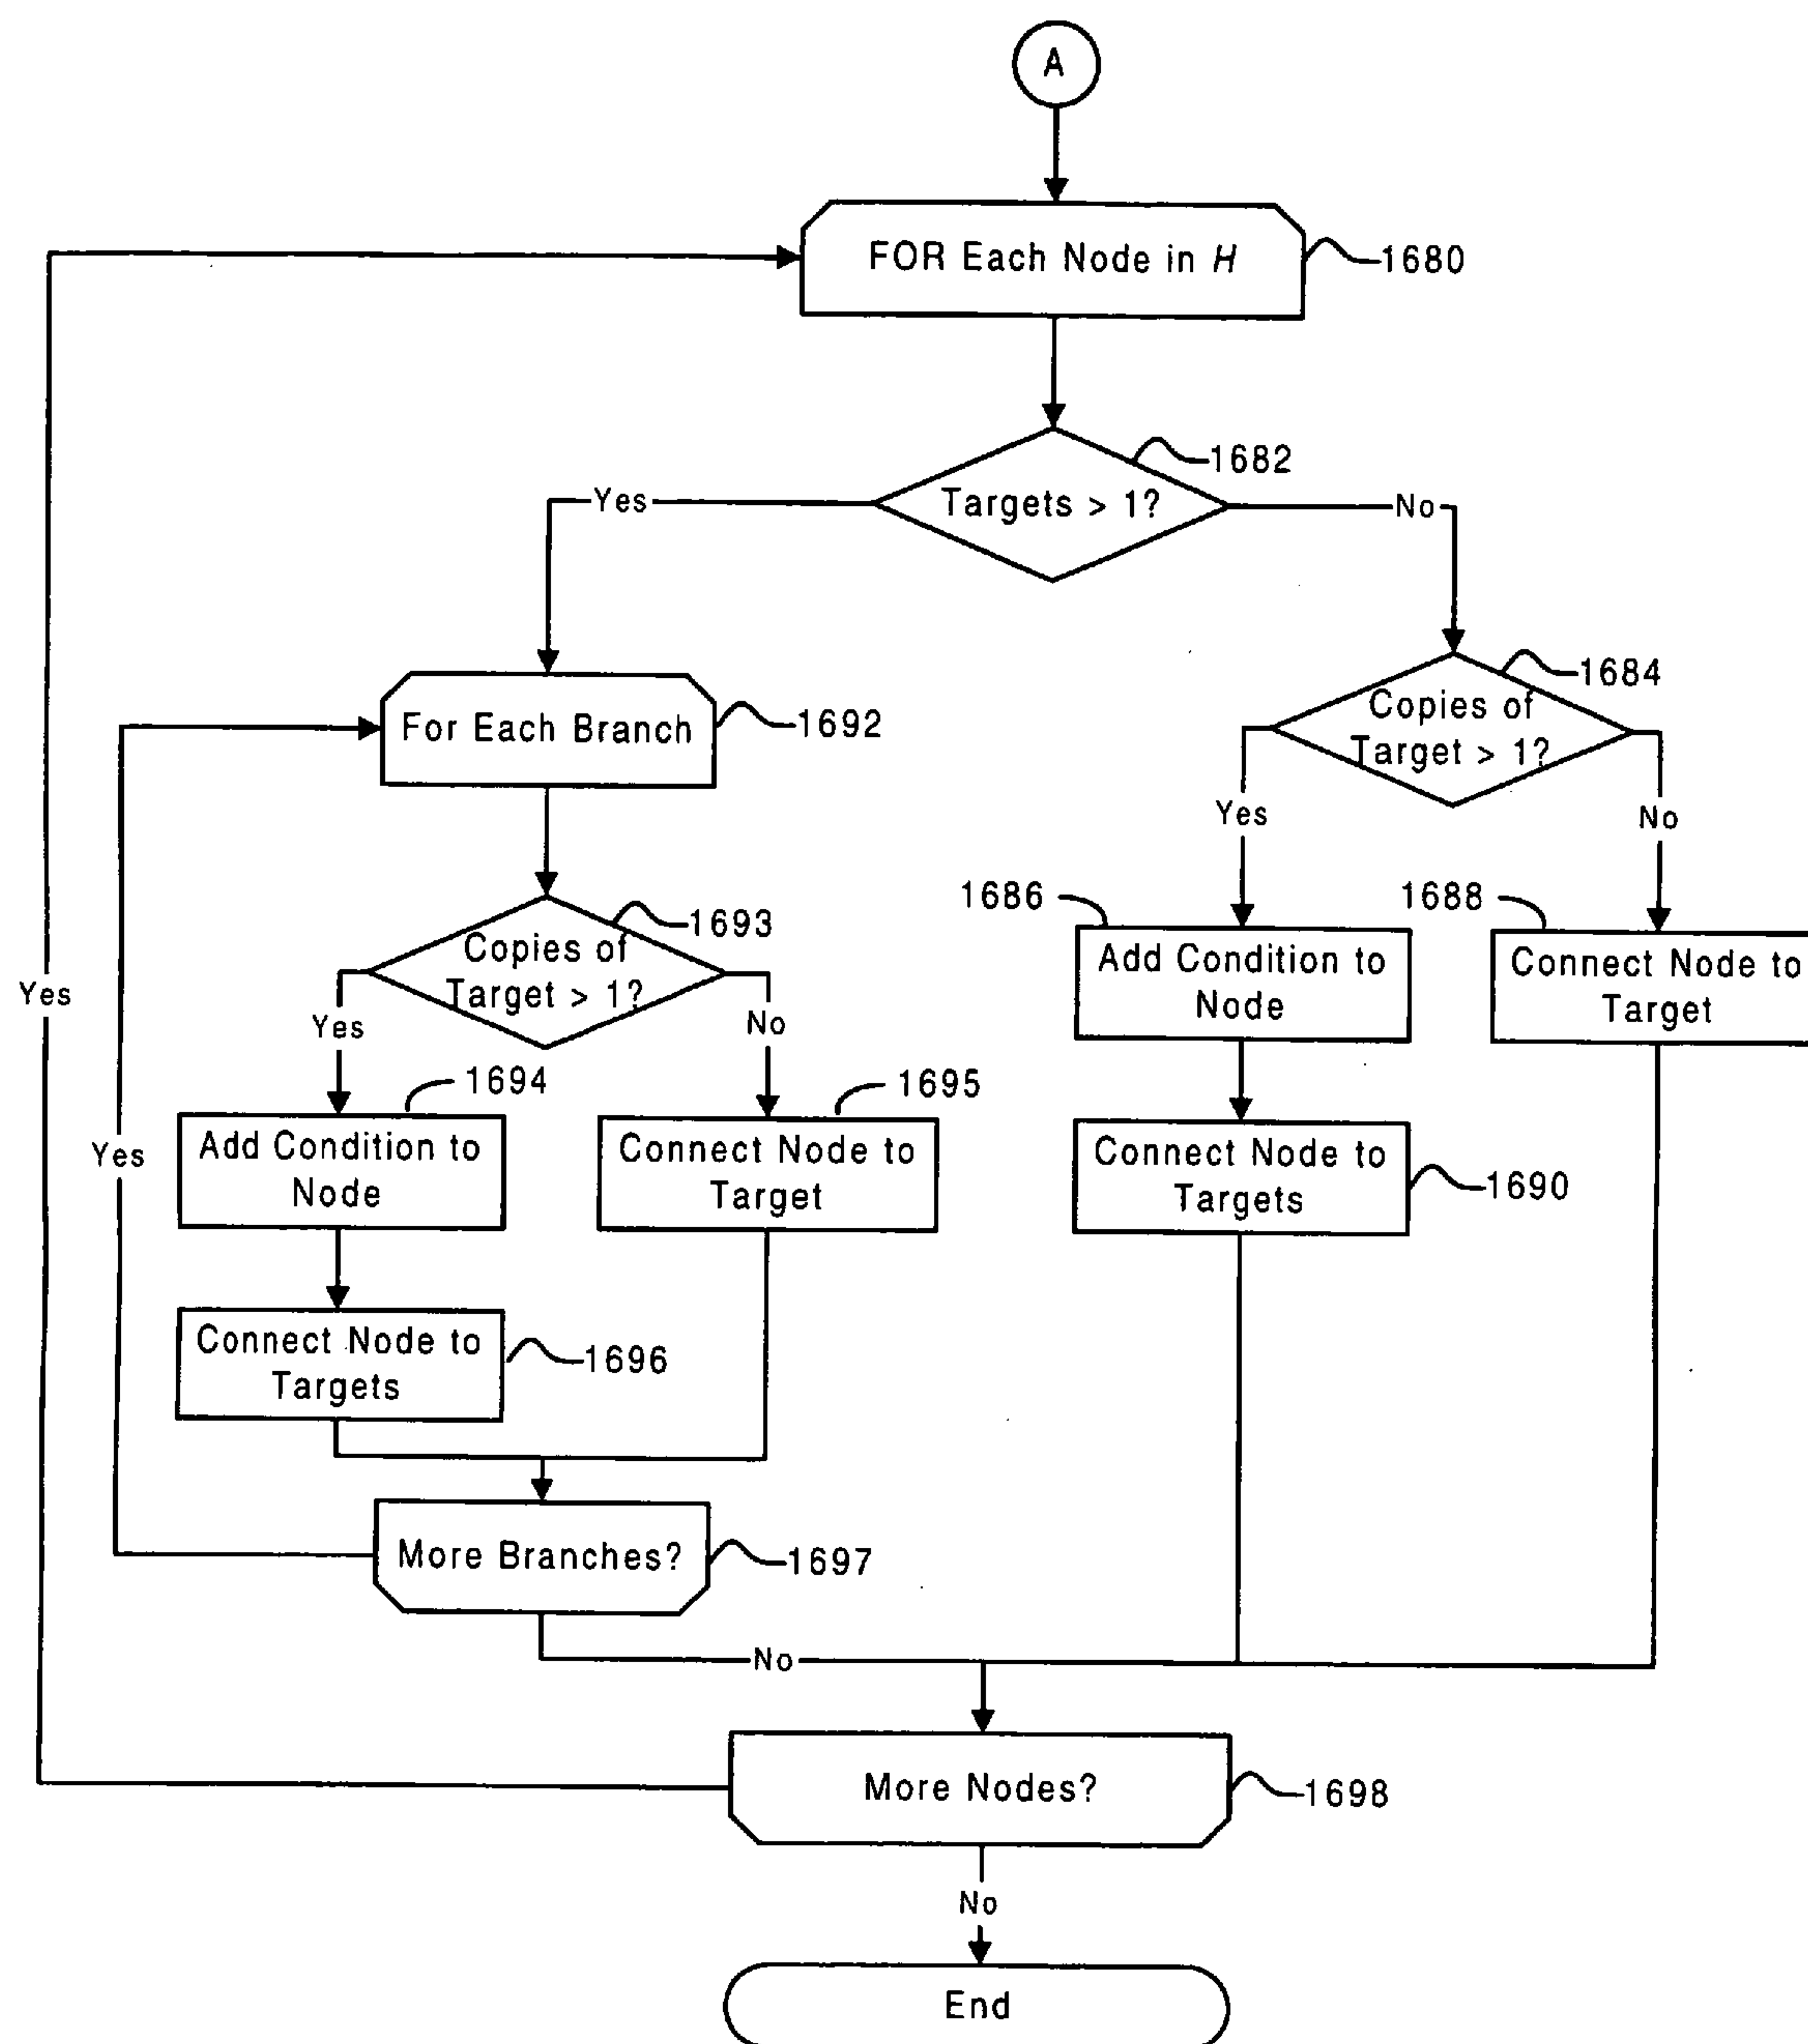


FIG. 16D




FIG. 17A




FIG. 17B

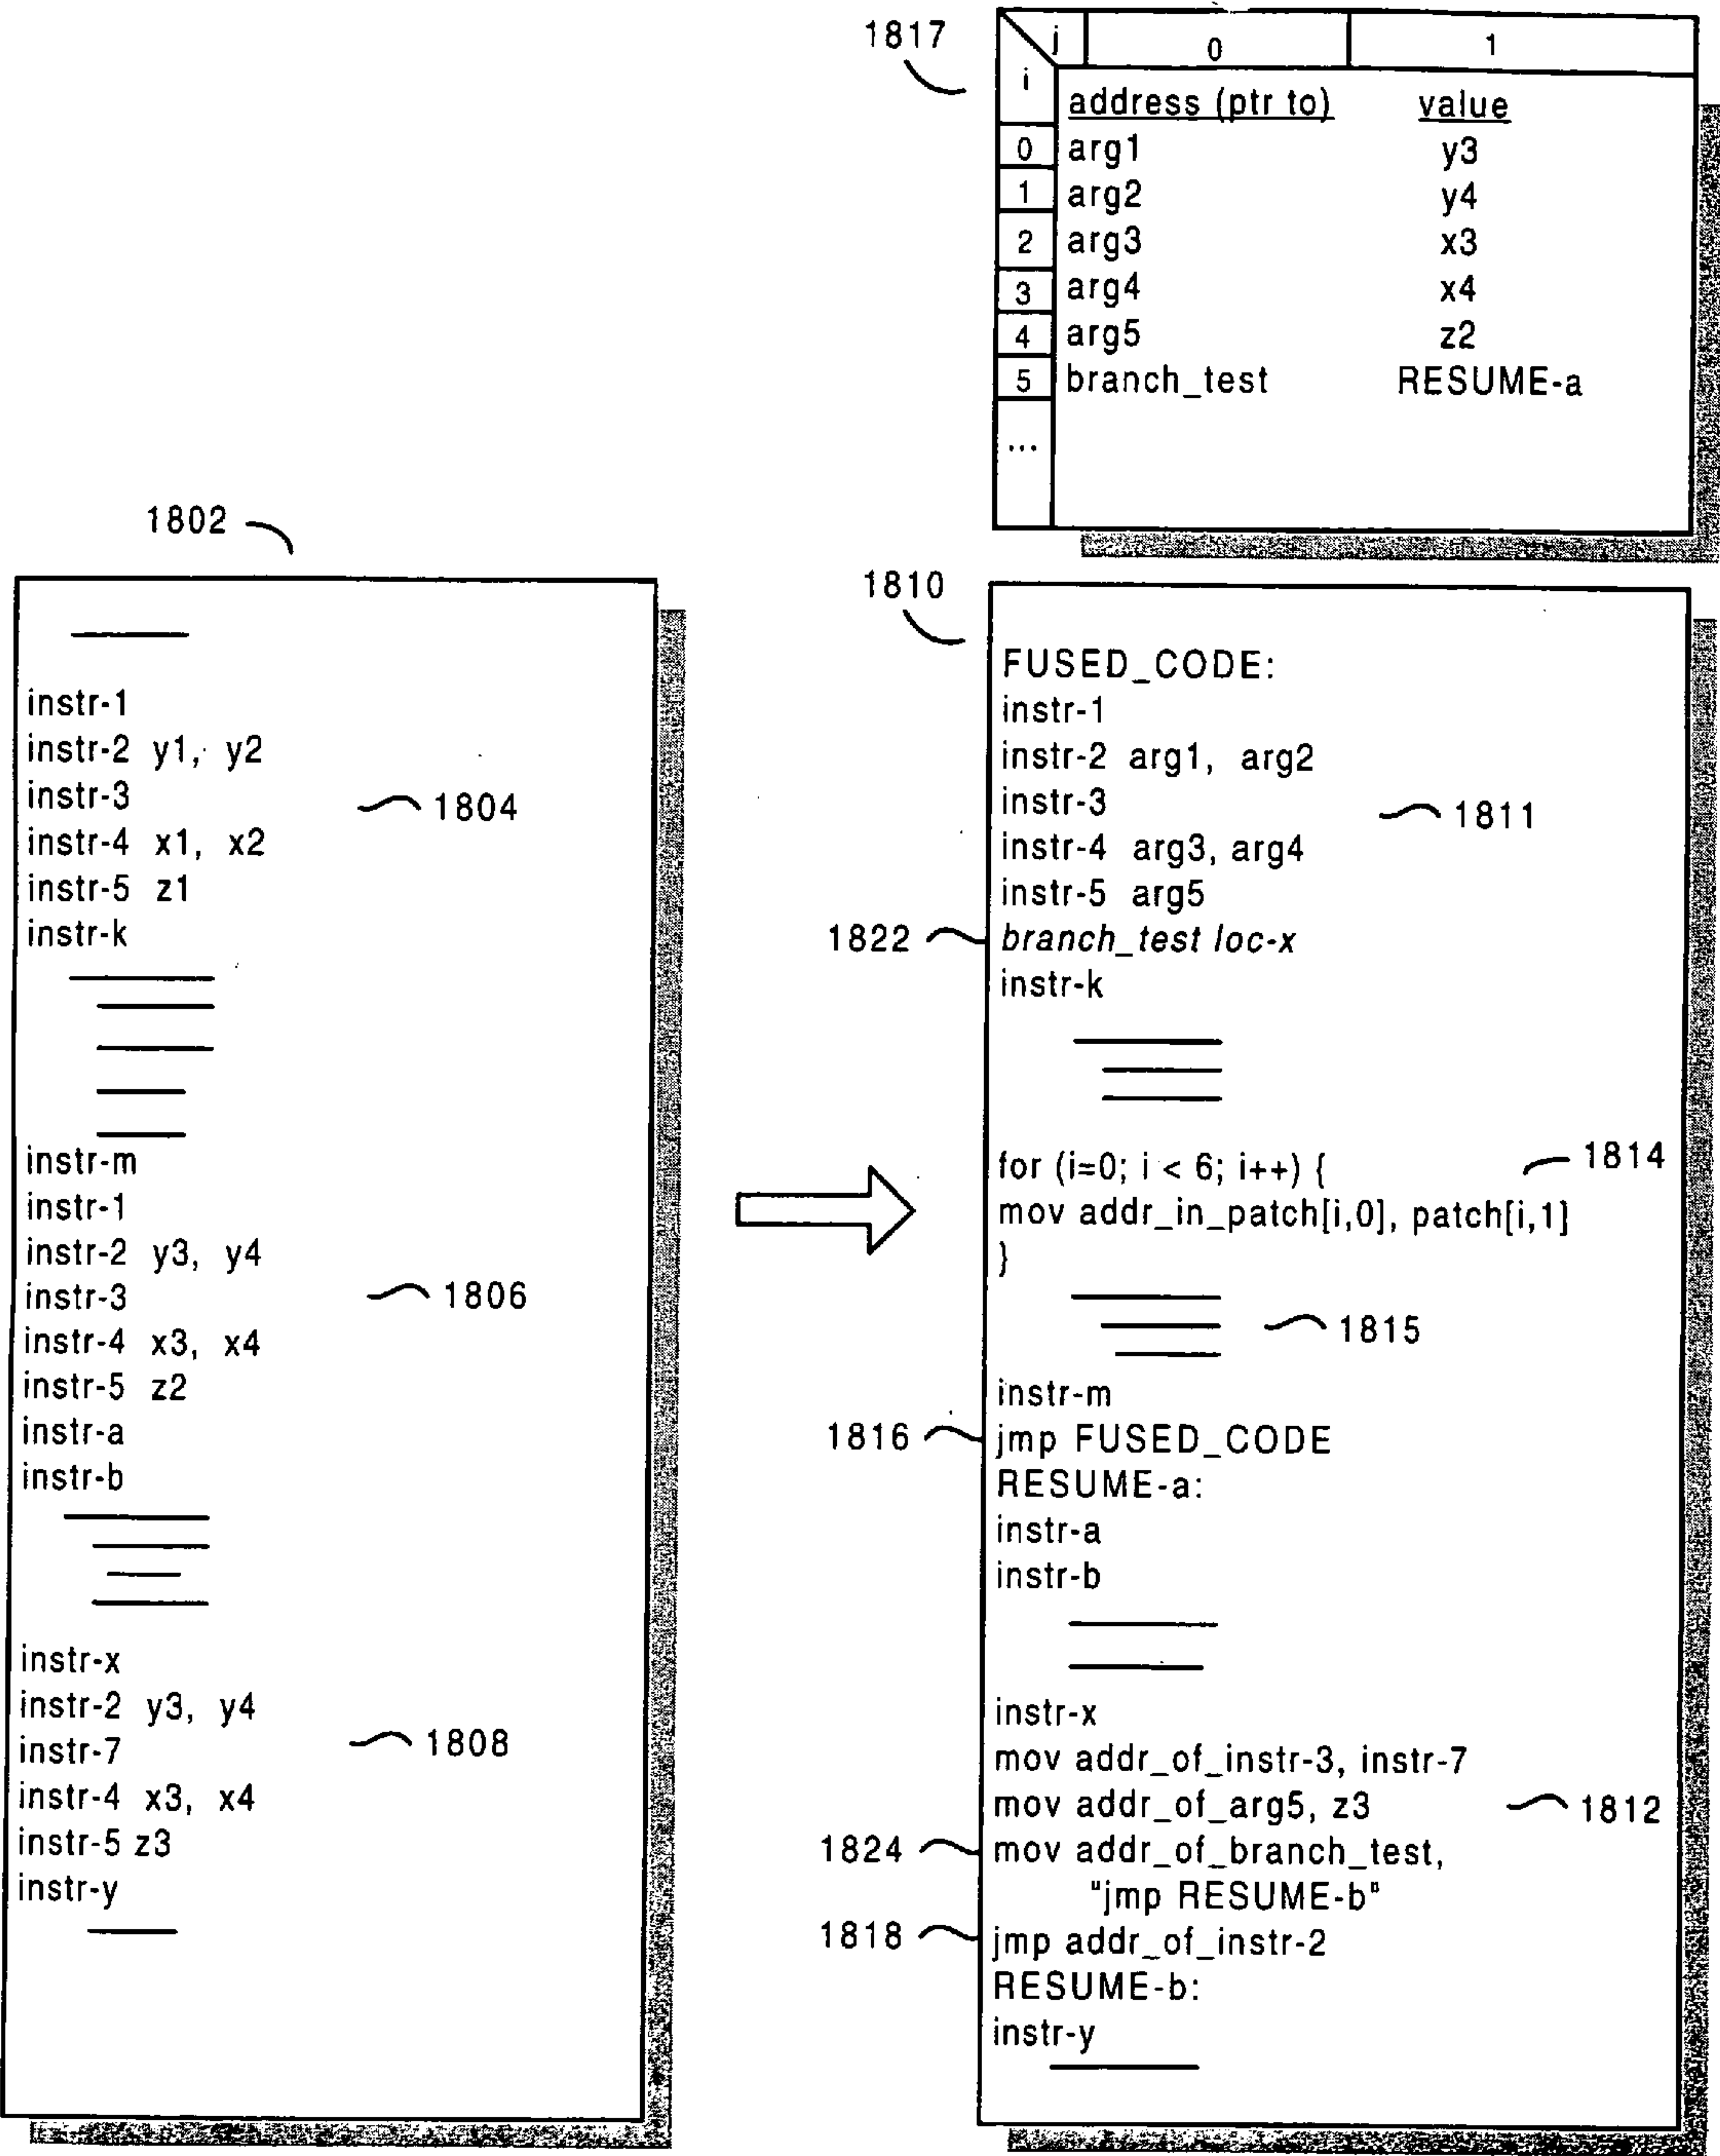


Fig. 18

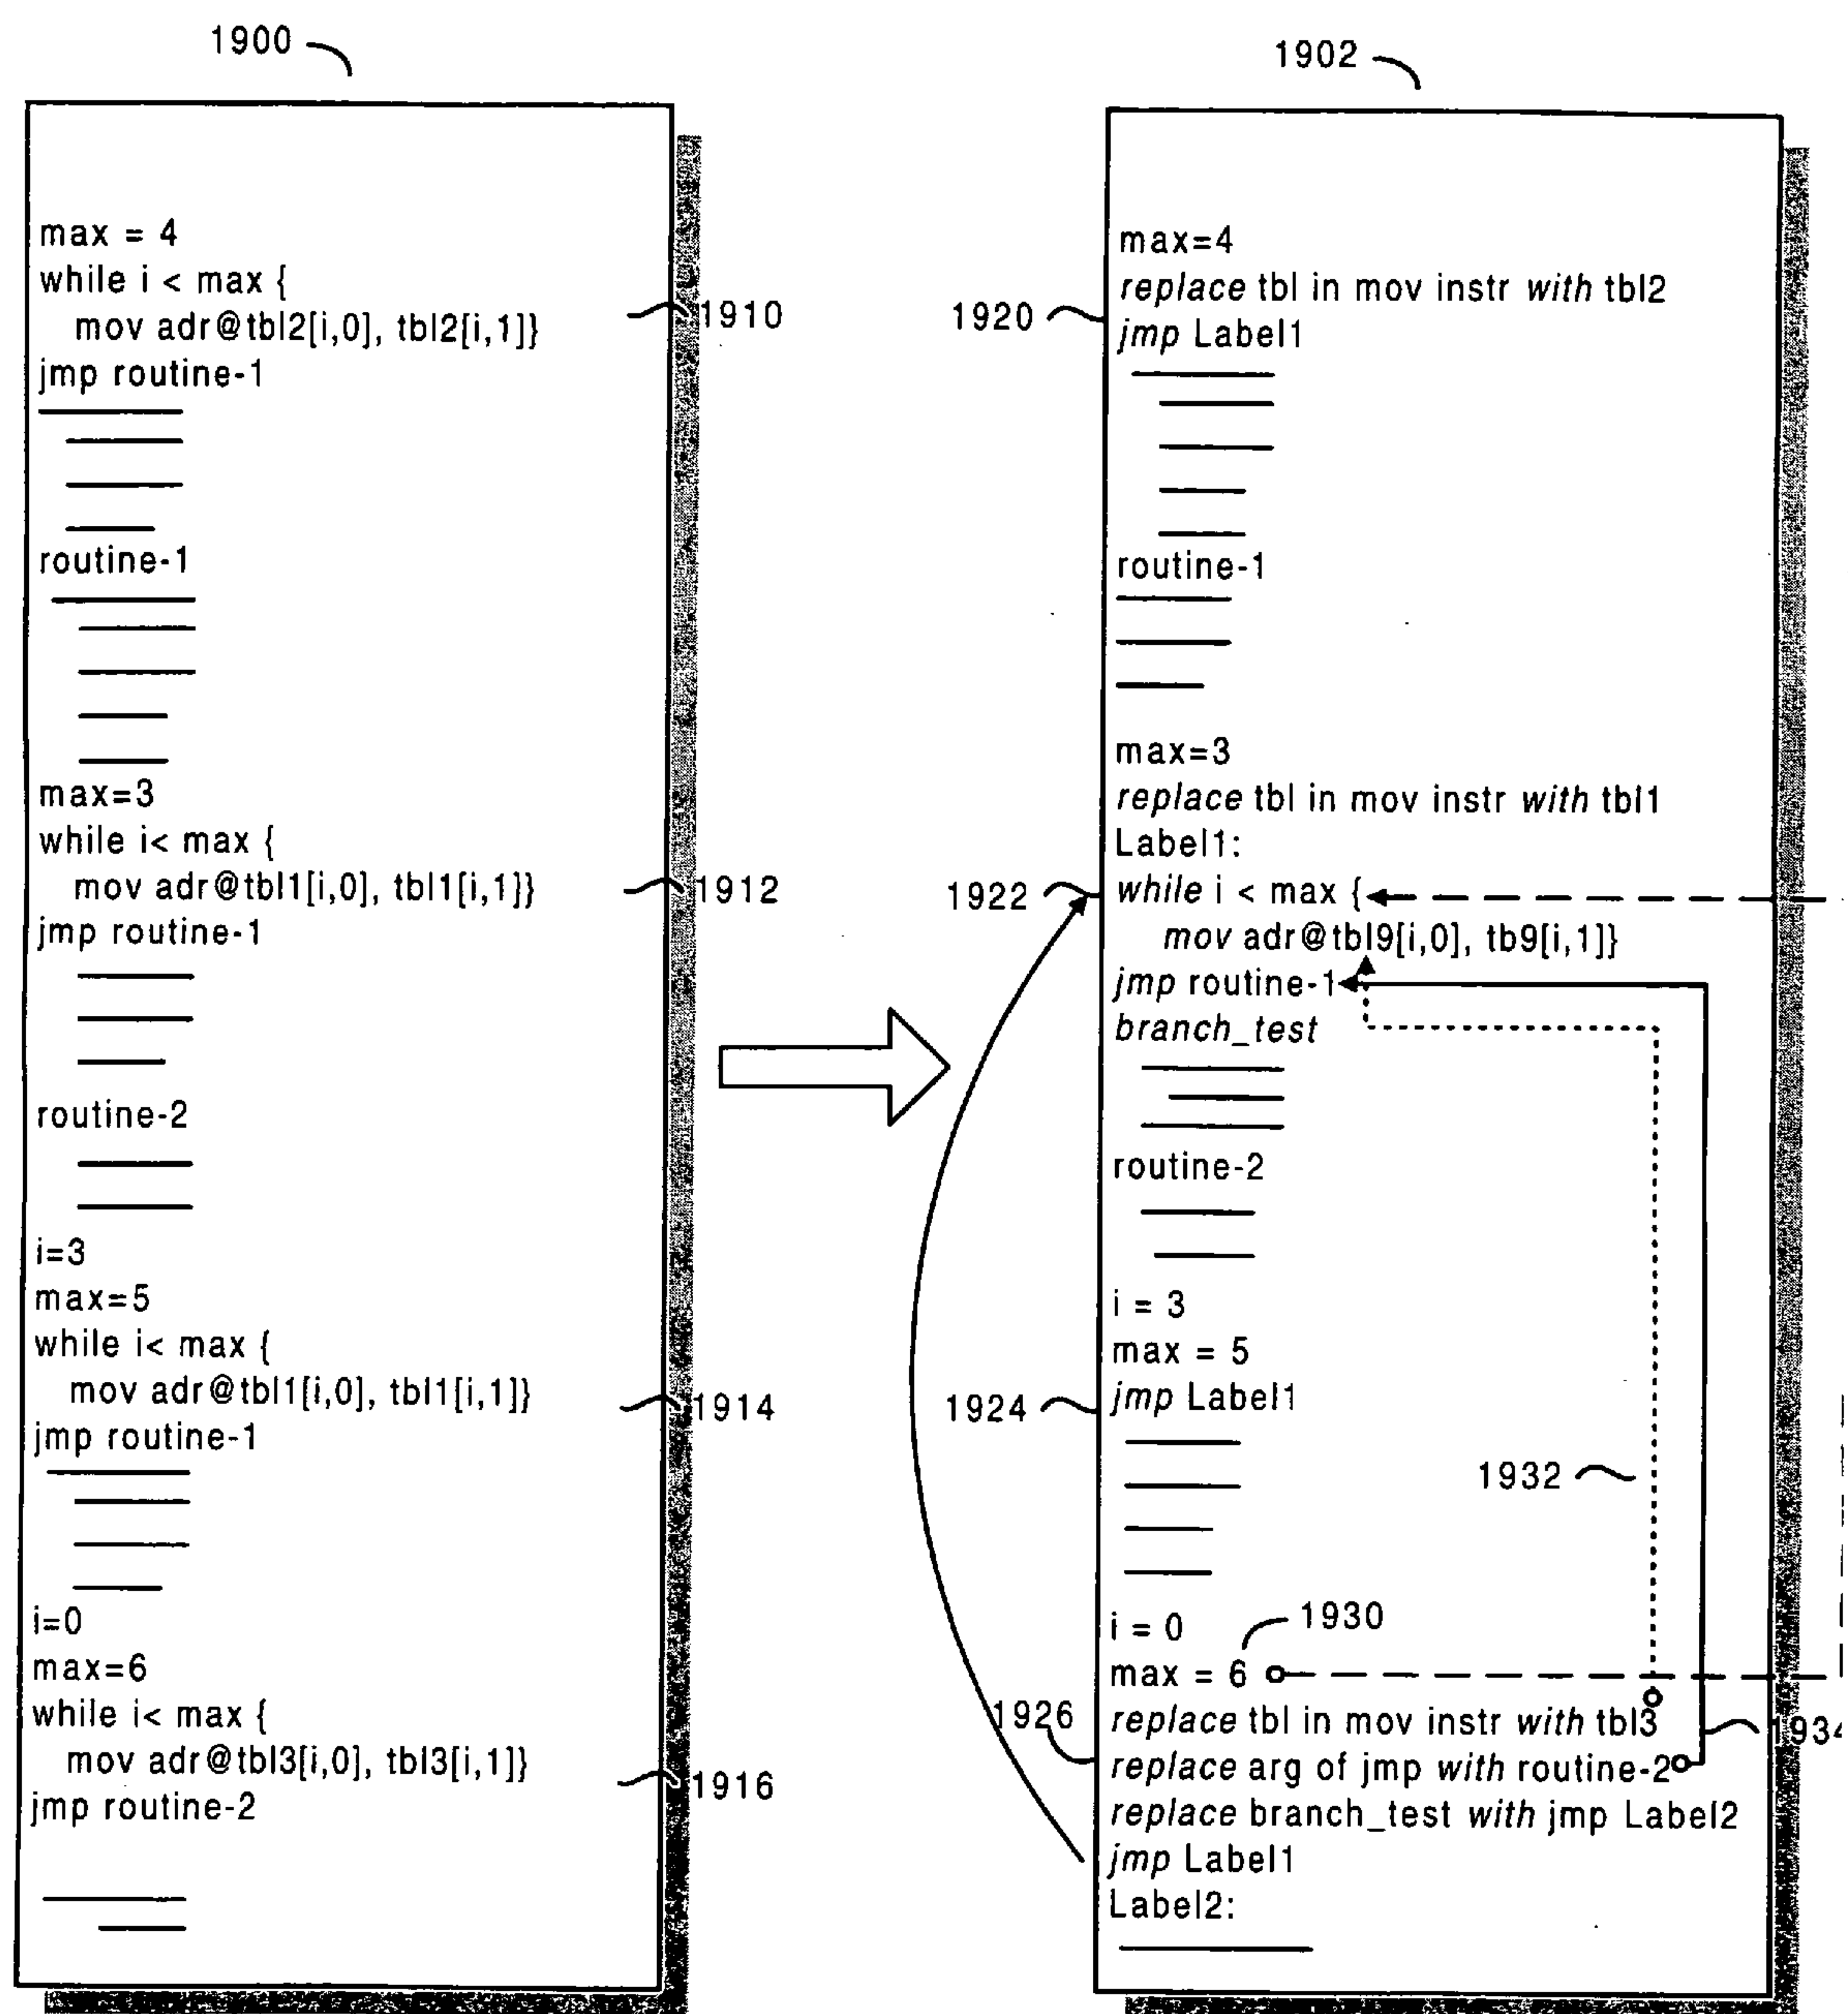


Fig. 19

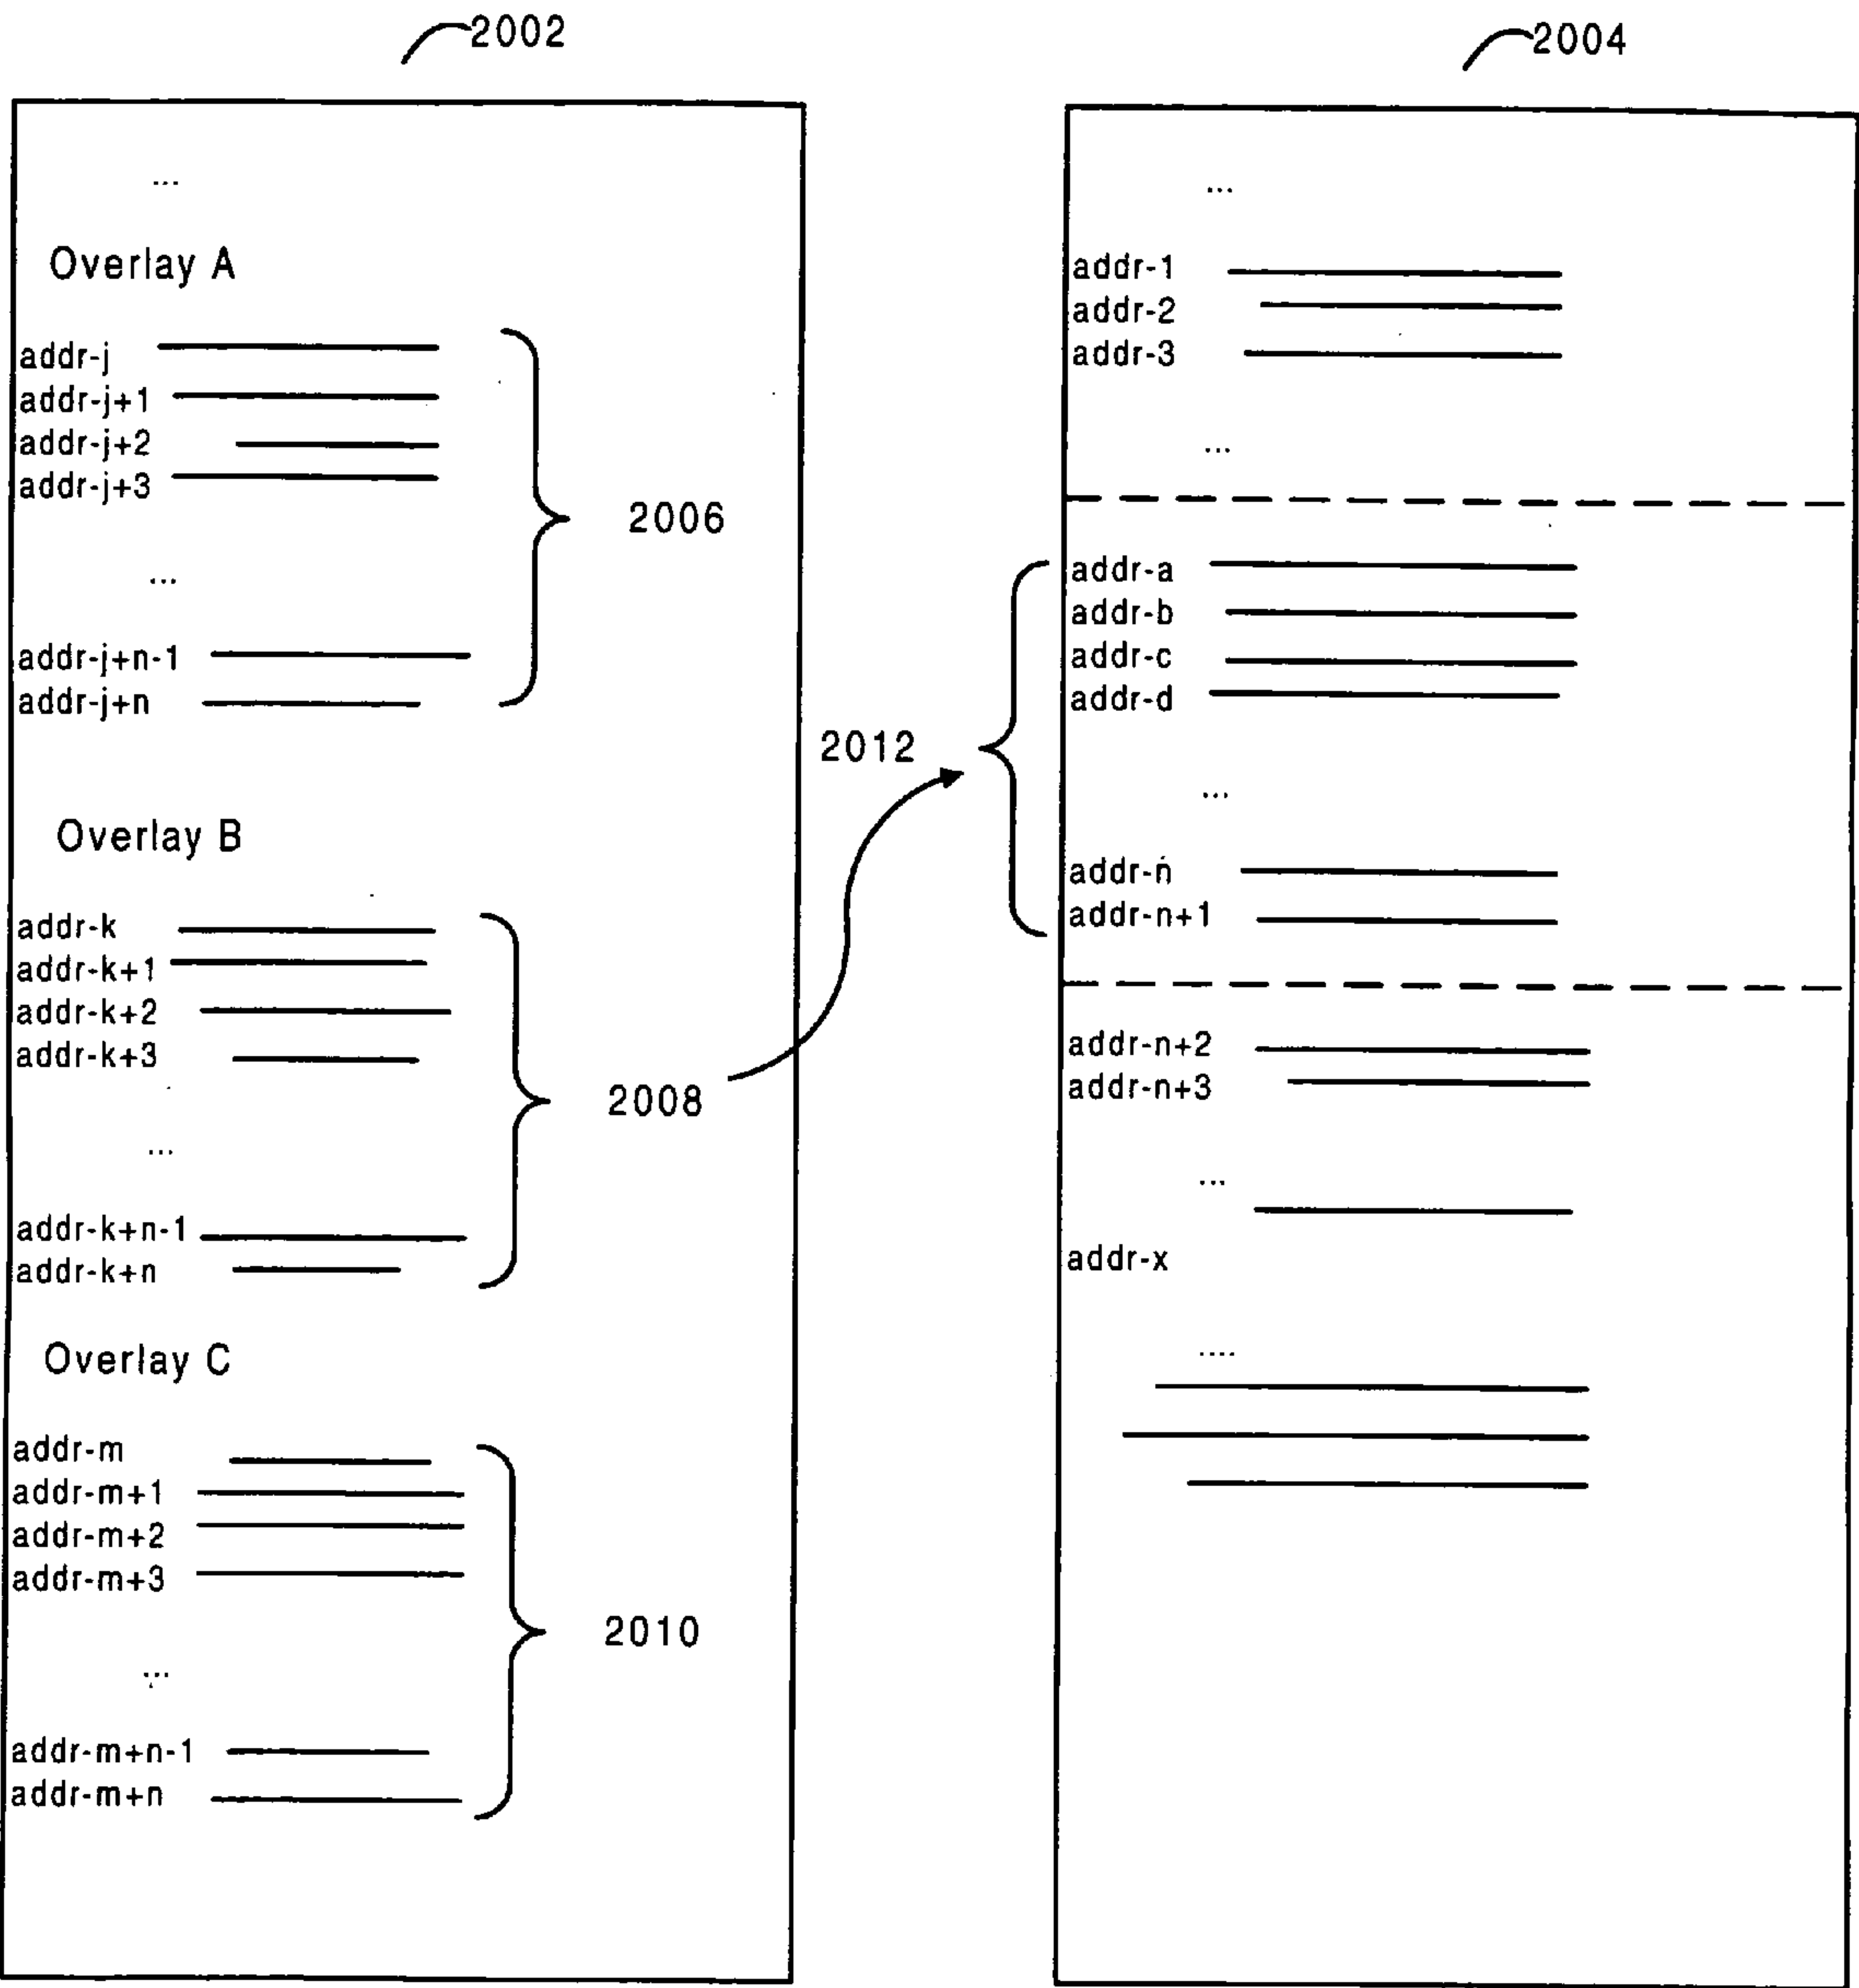


Fig. 20




Fig. 21

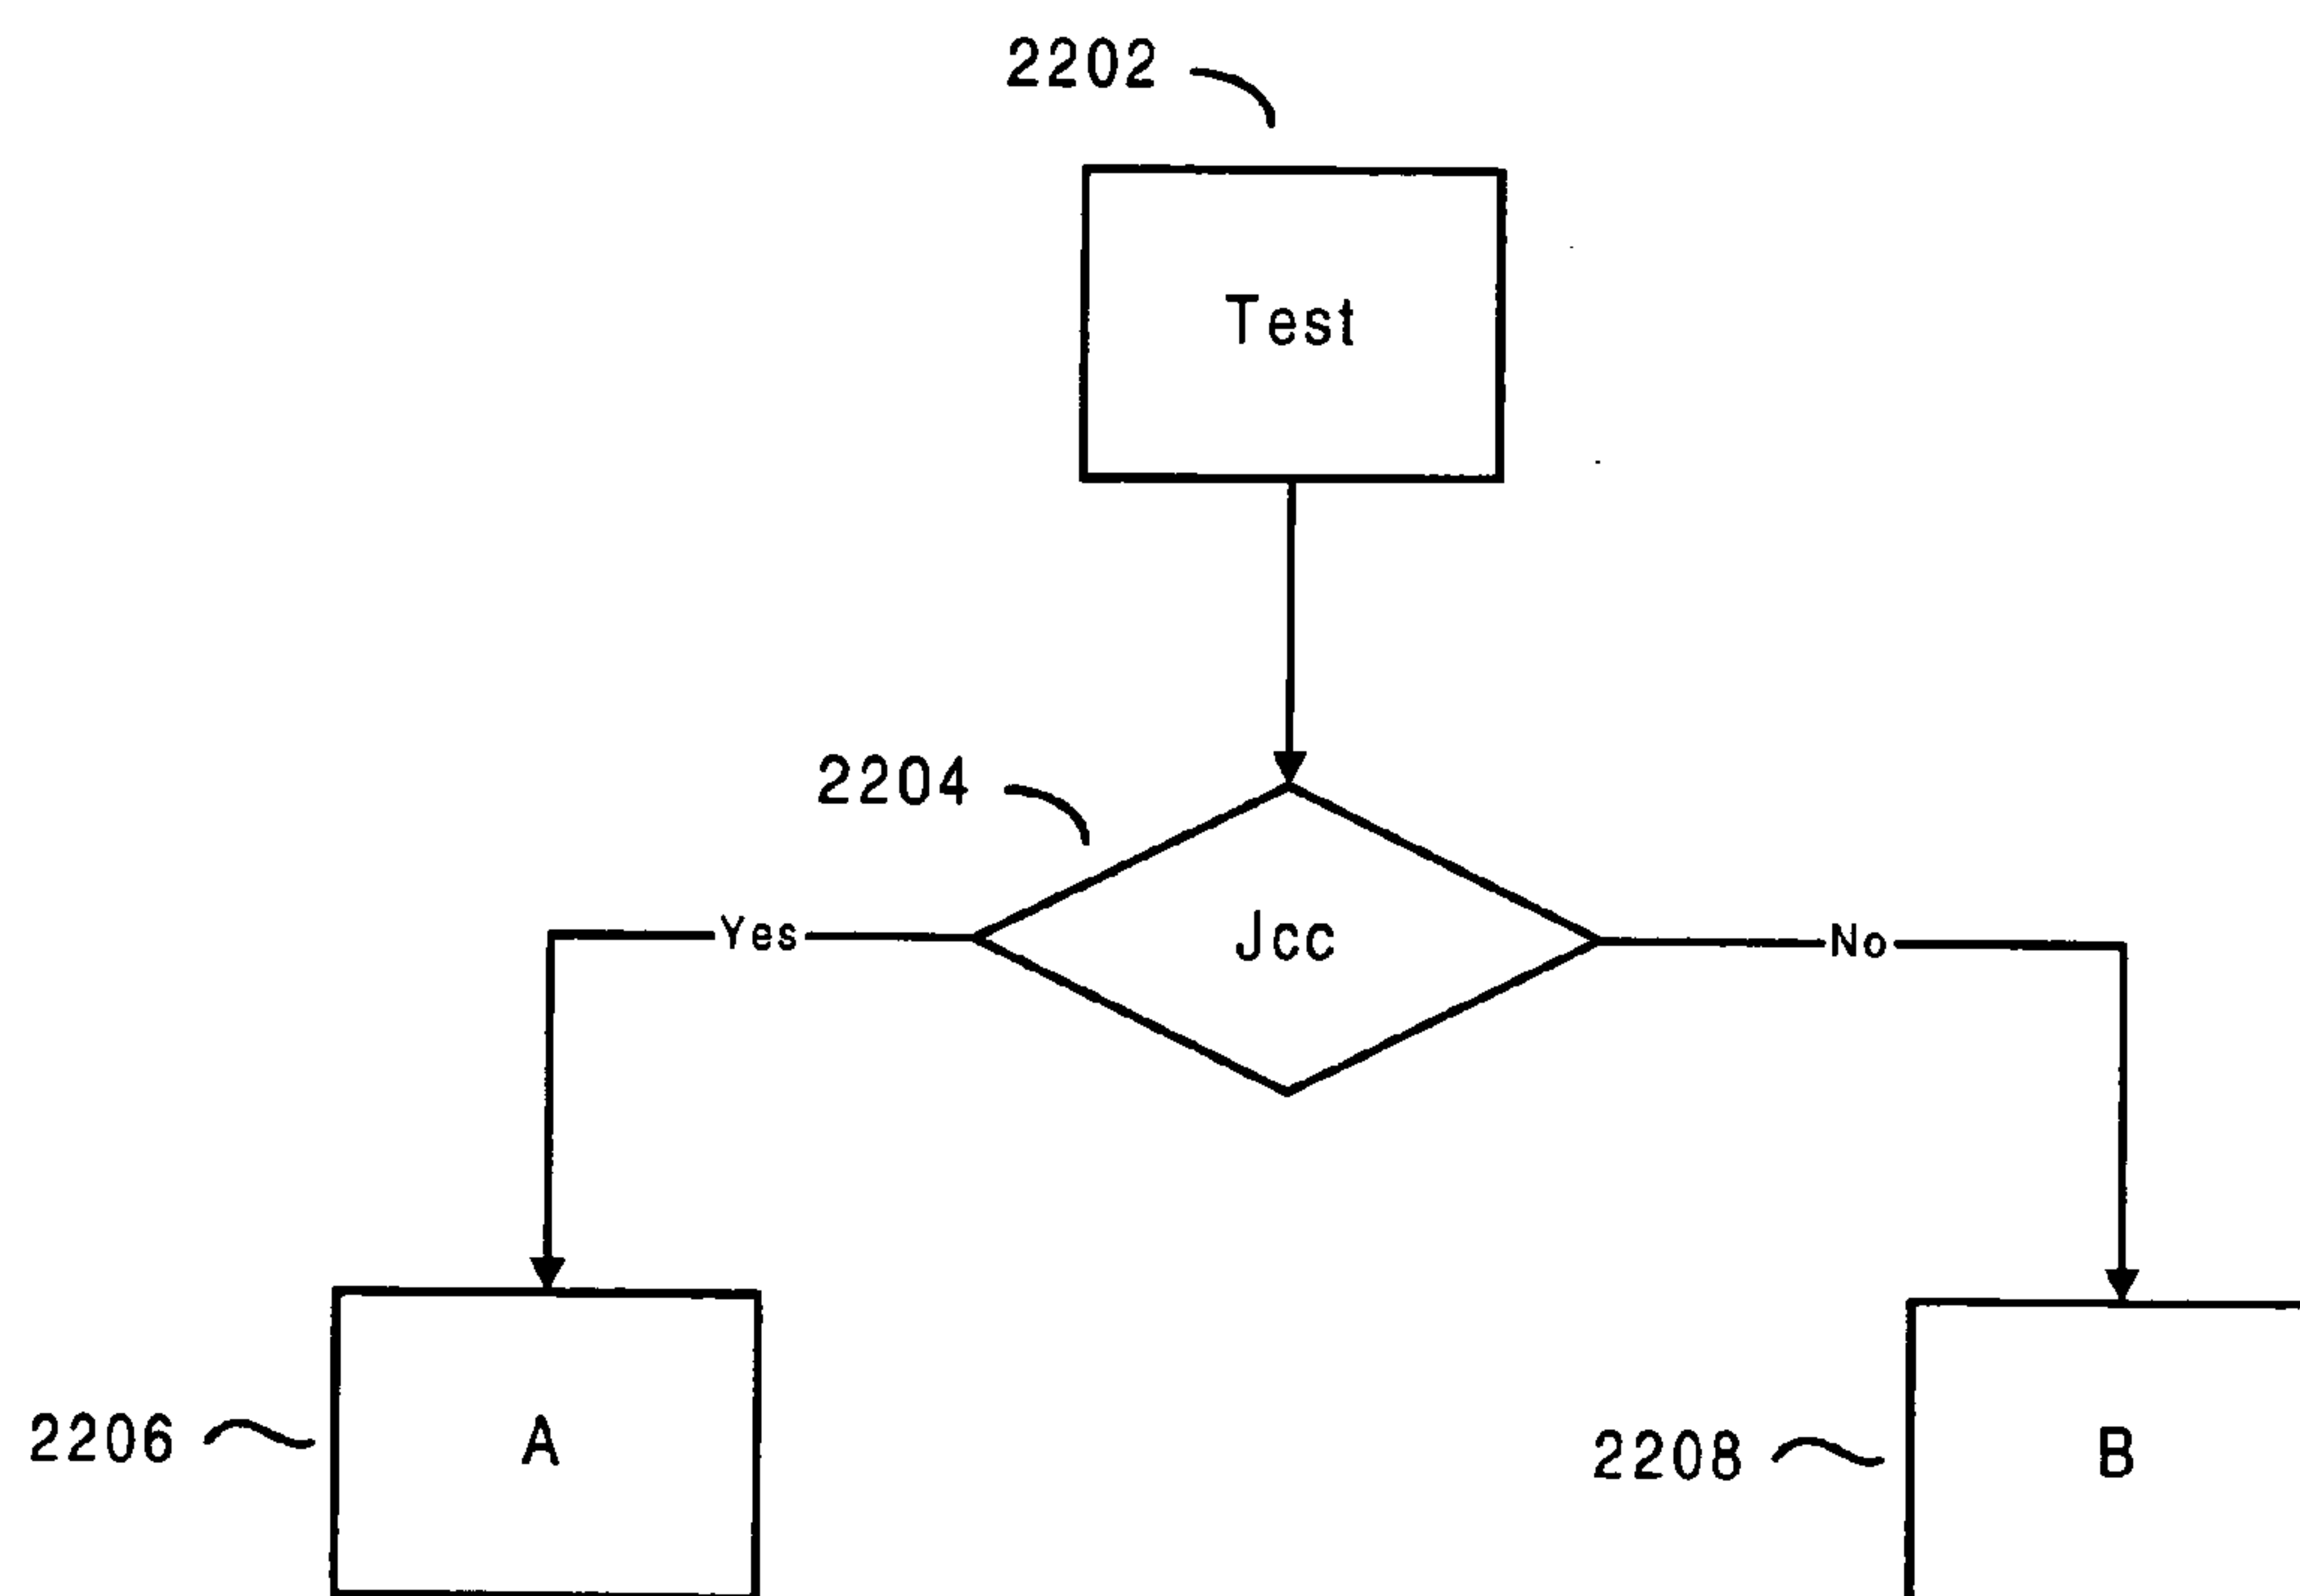


Fig. 22A

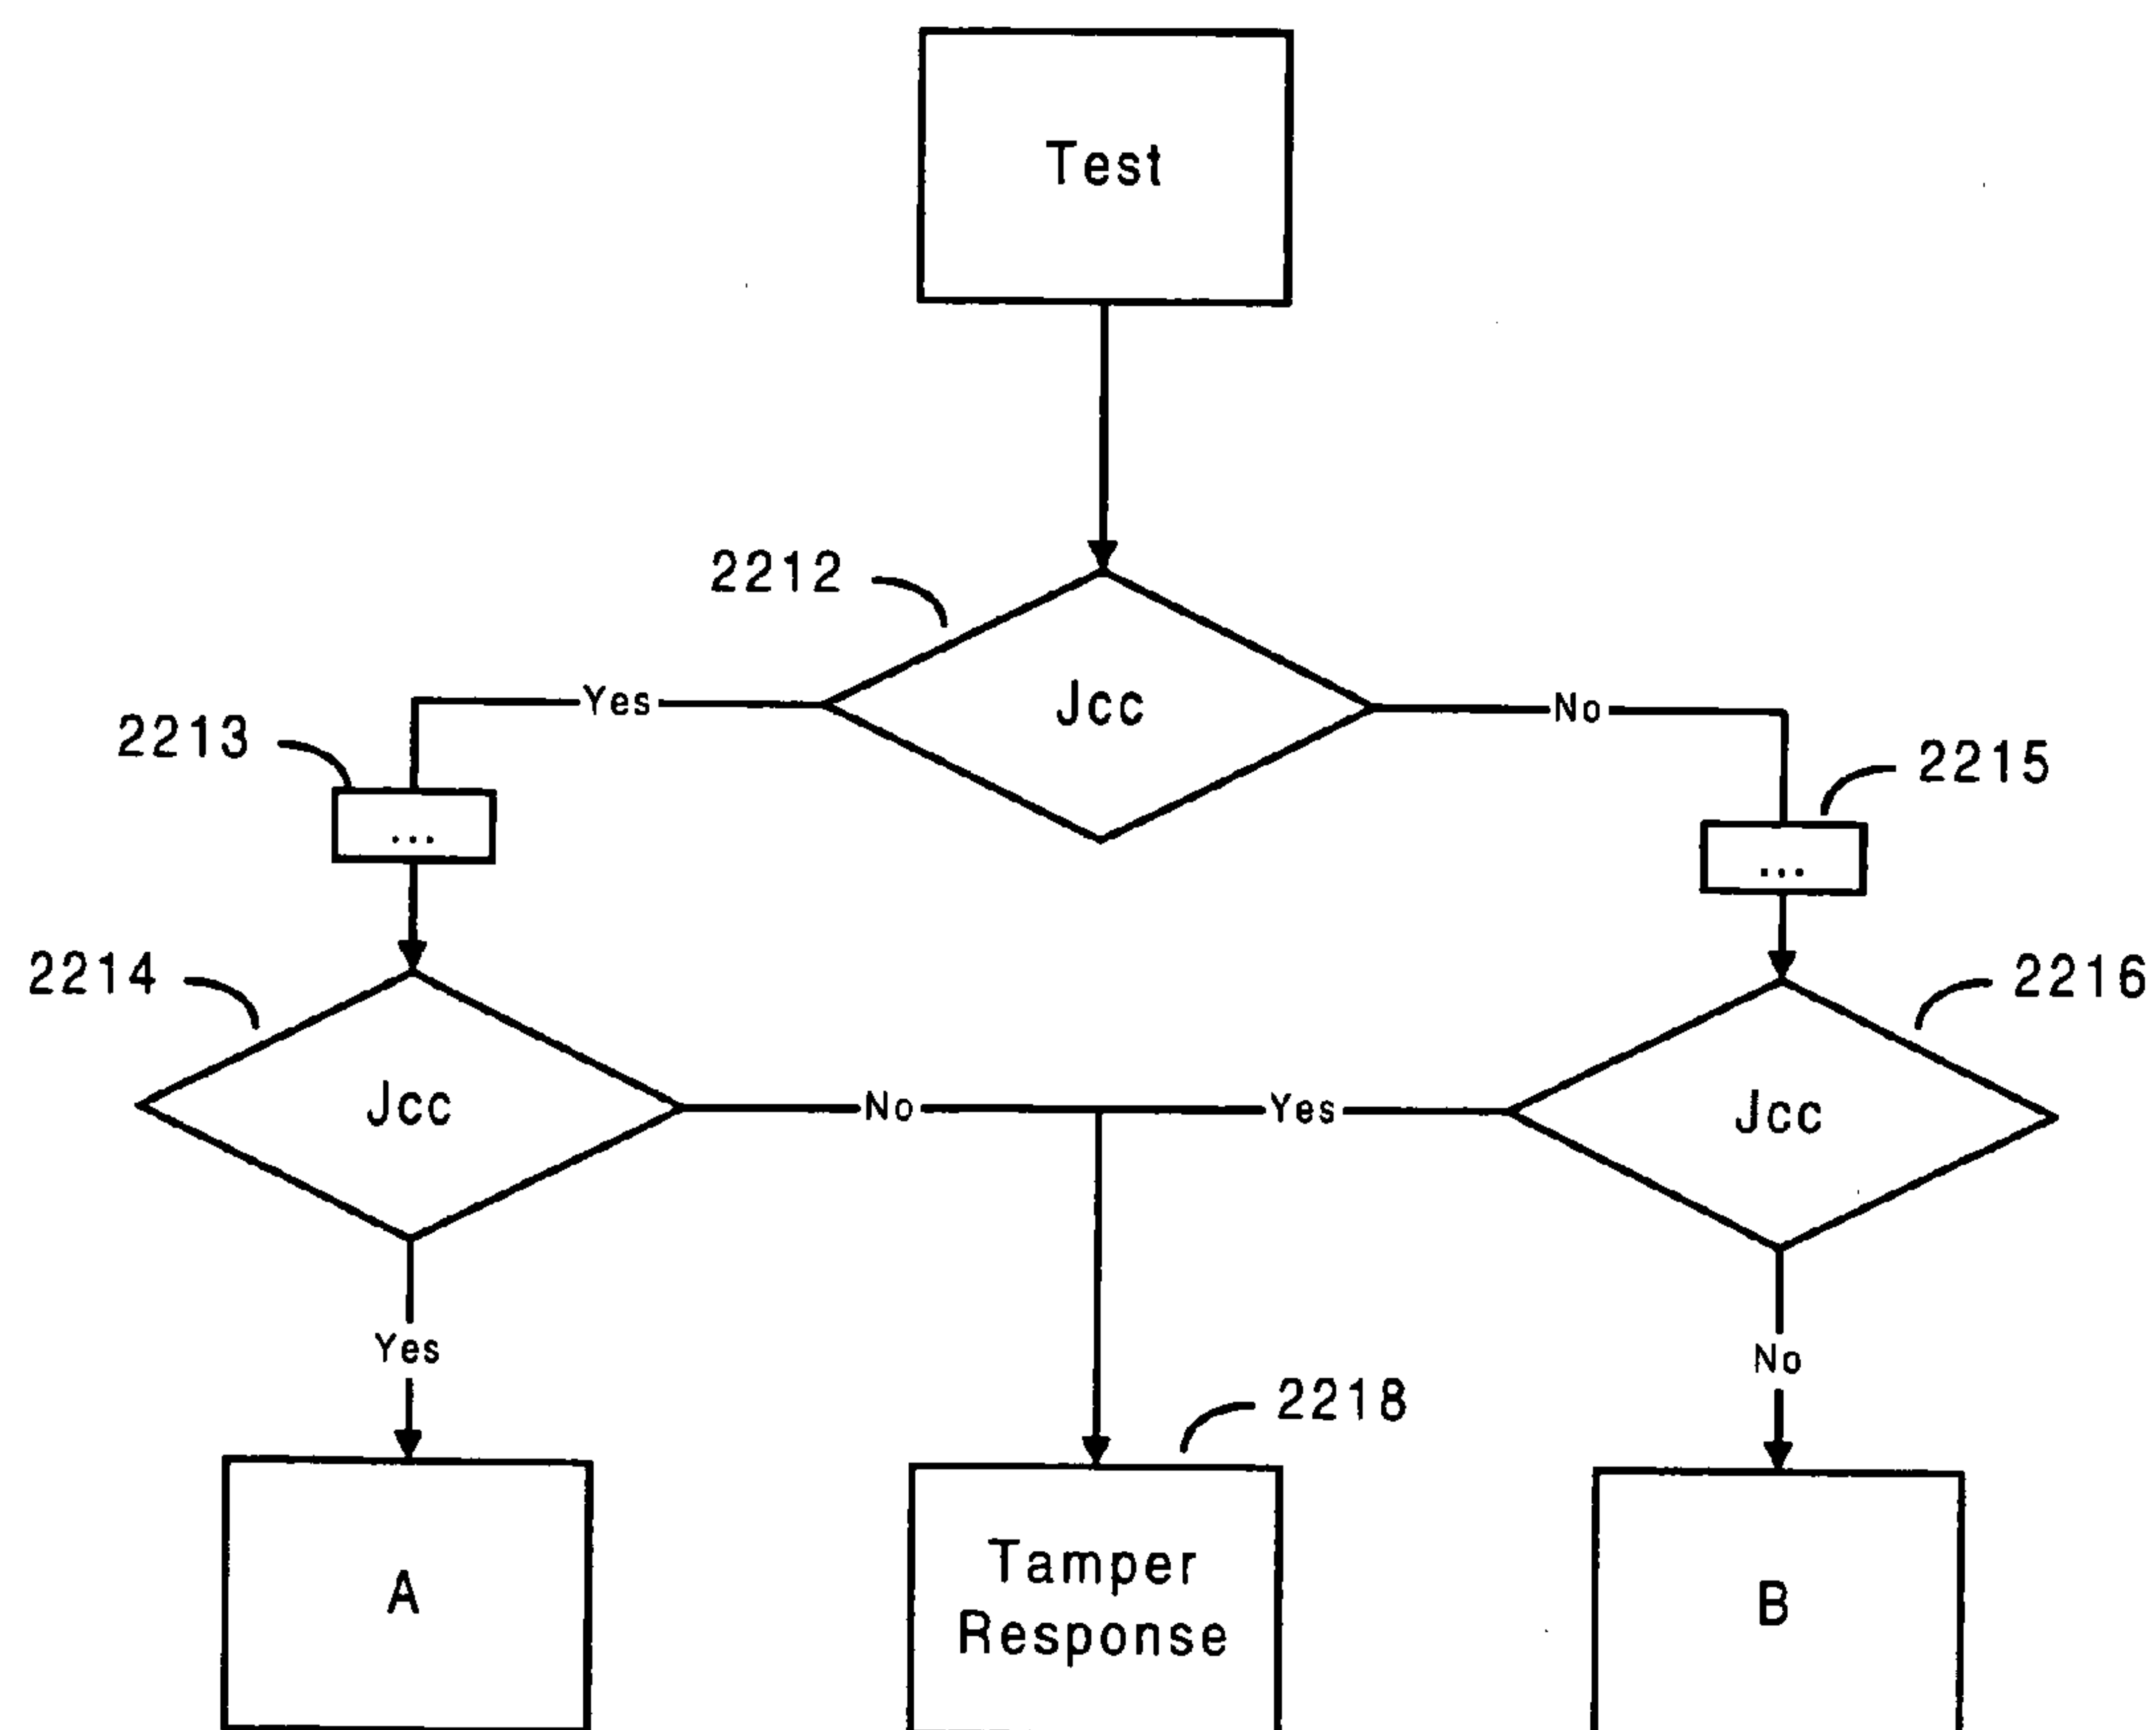


Fig. 22B

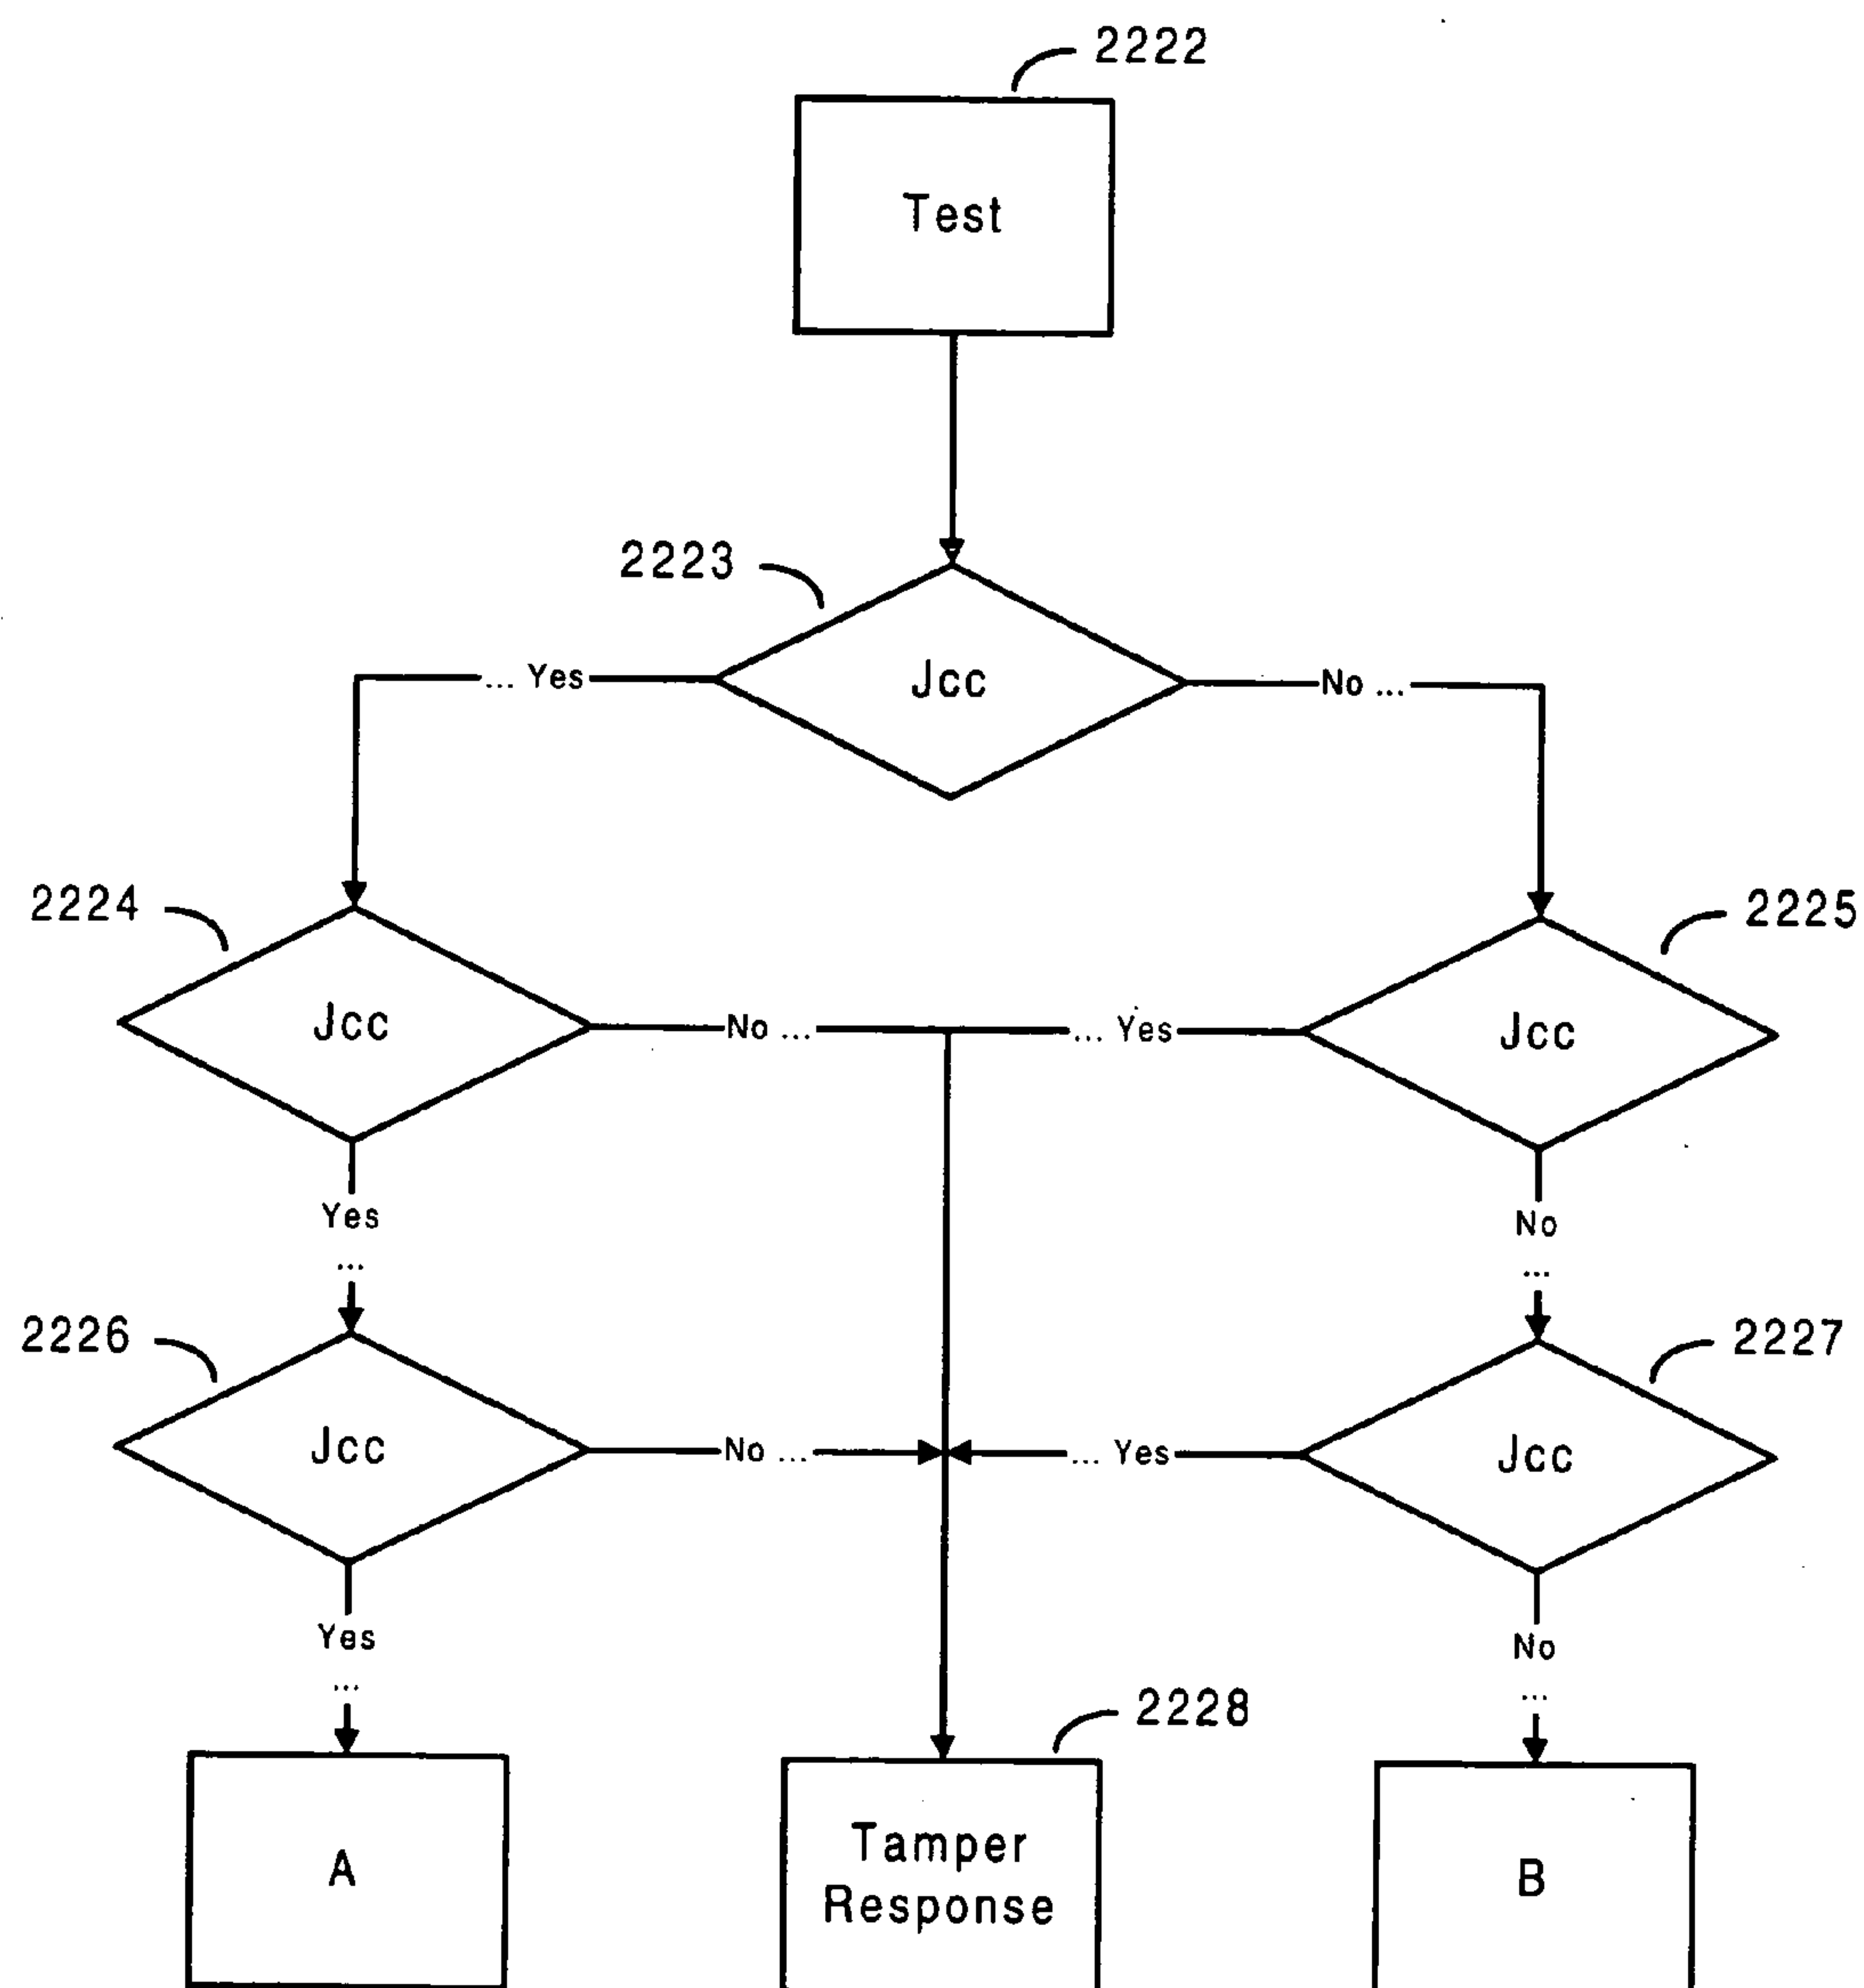


Fig. 22C

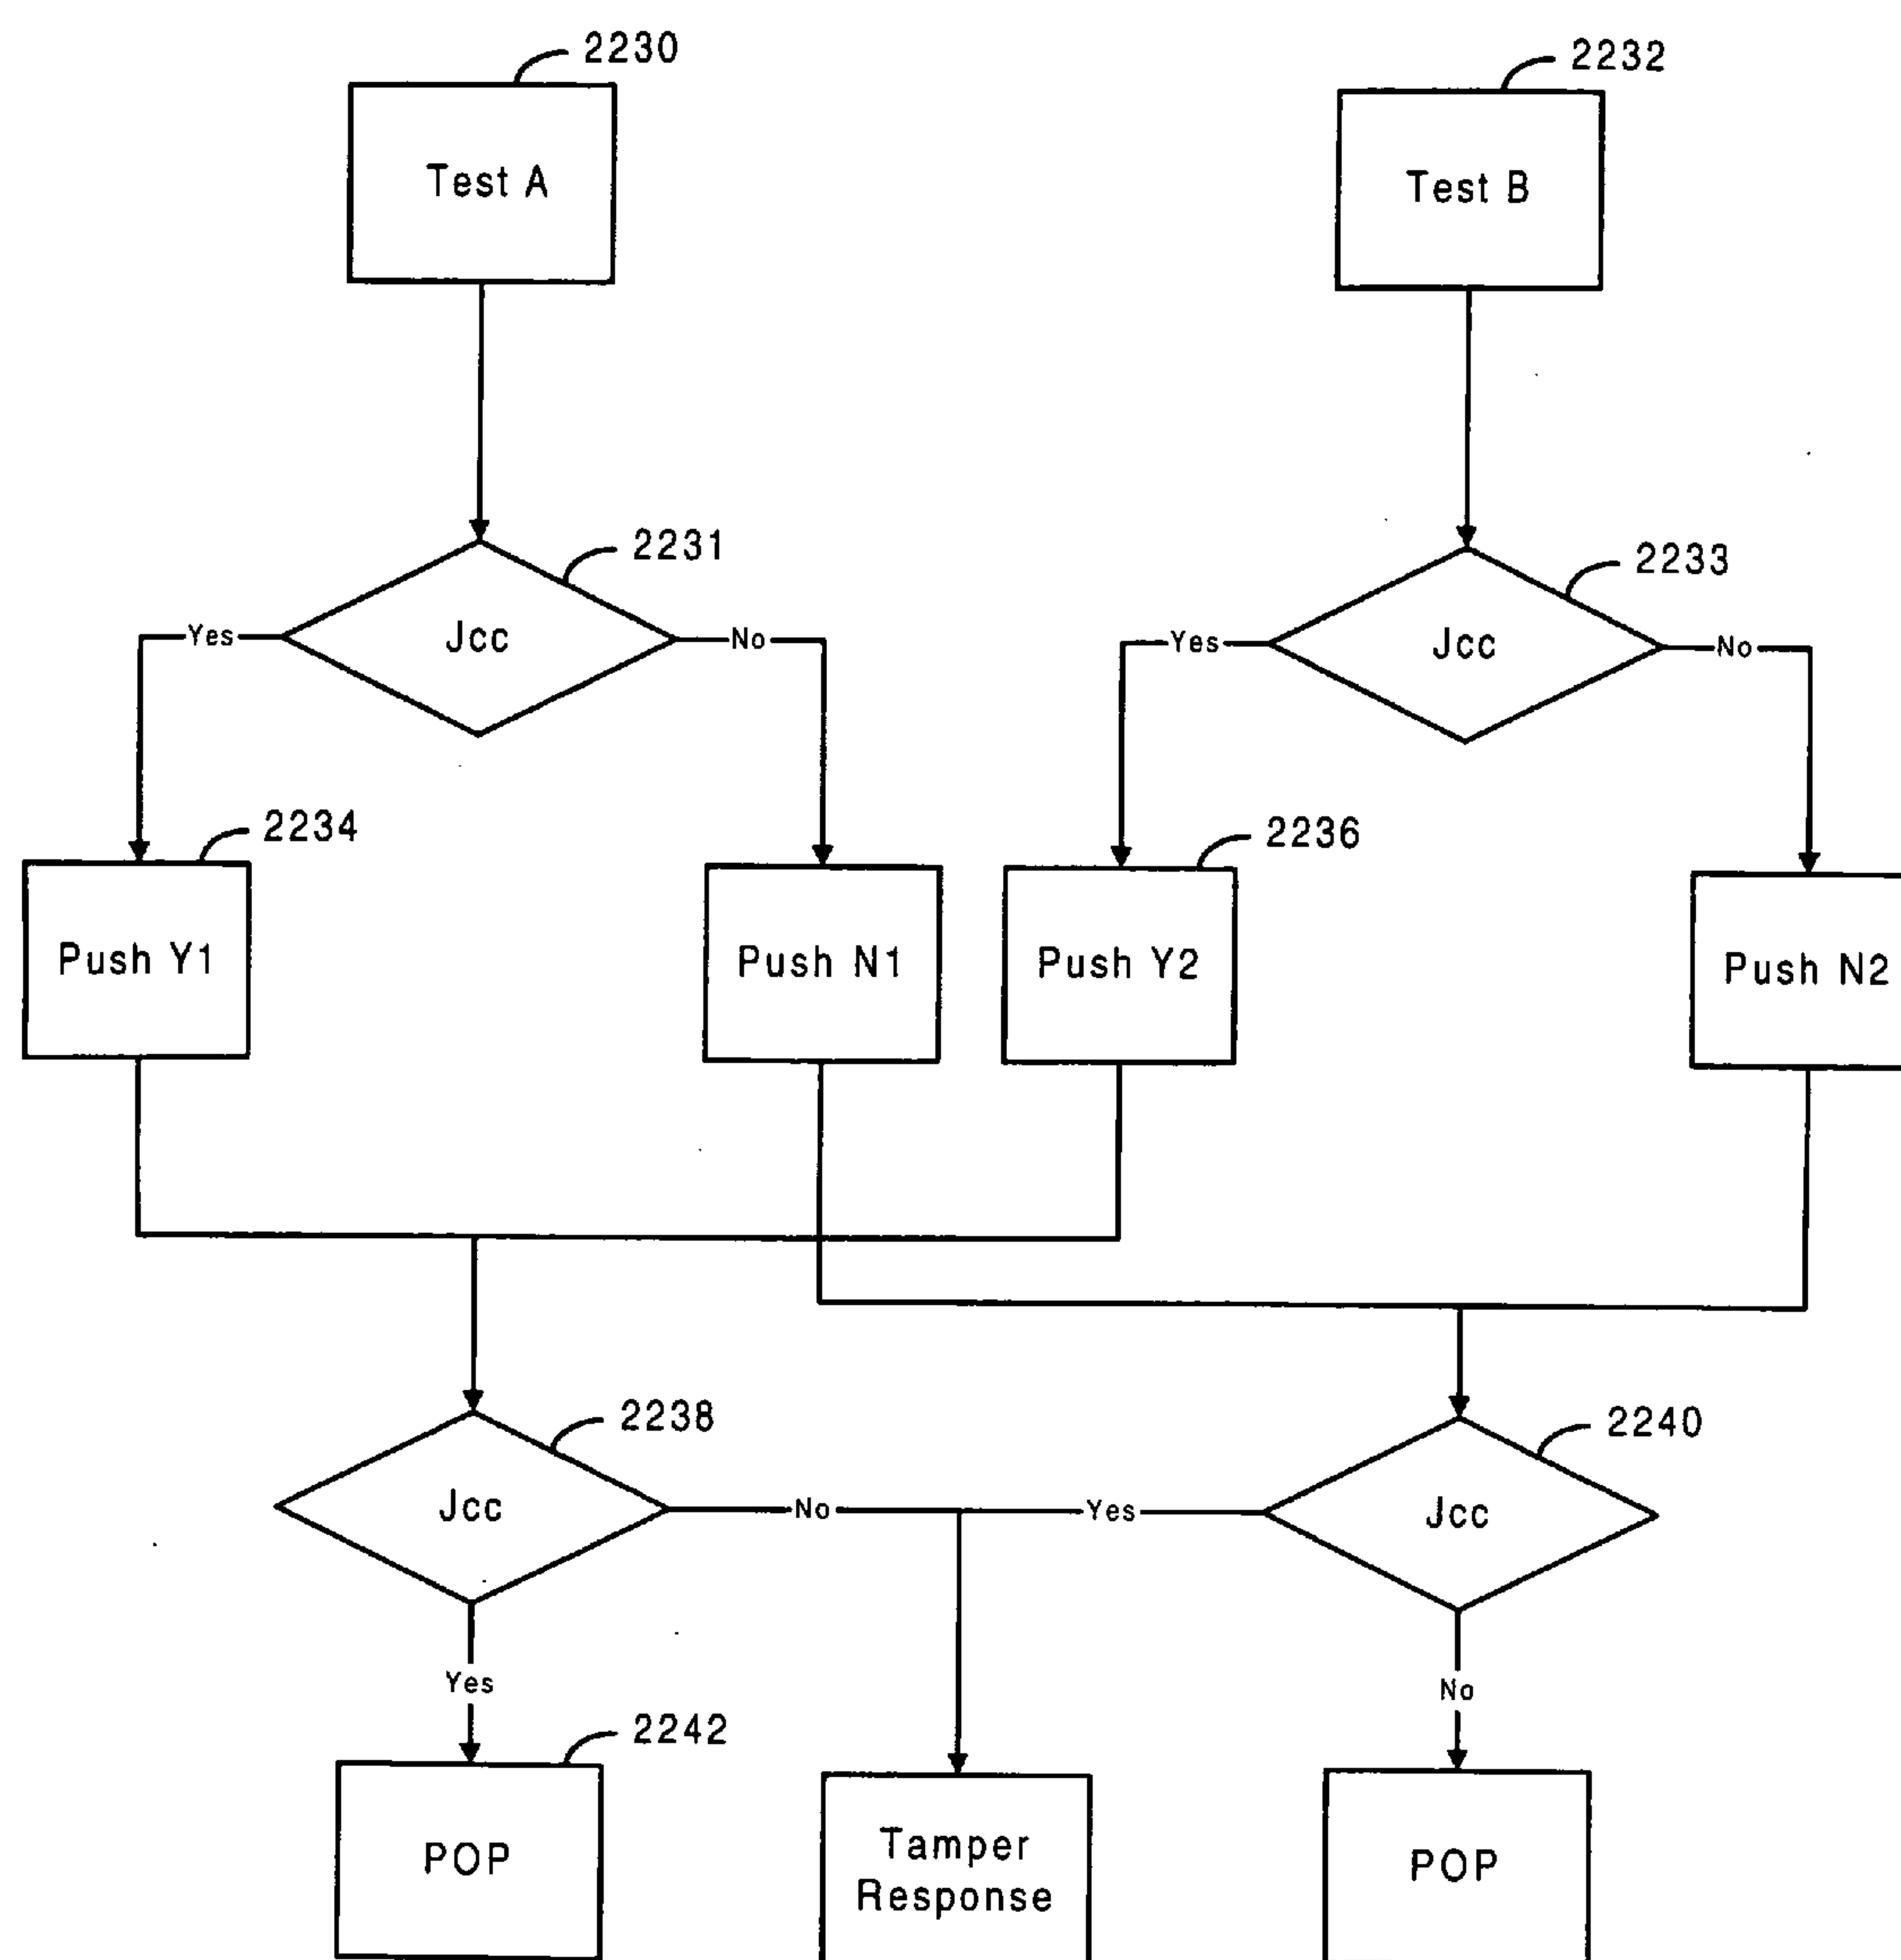


Fig. 22D

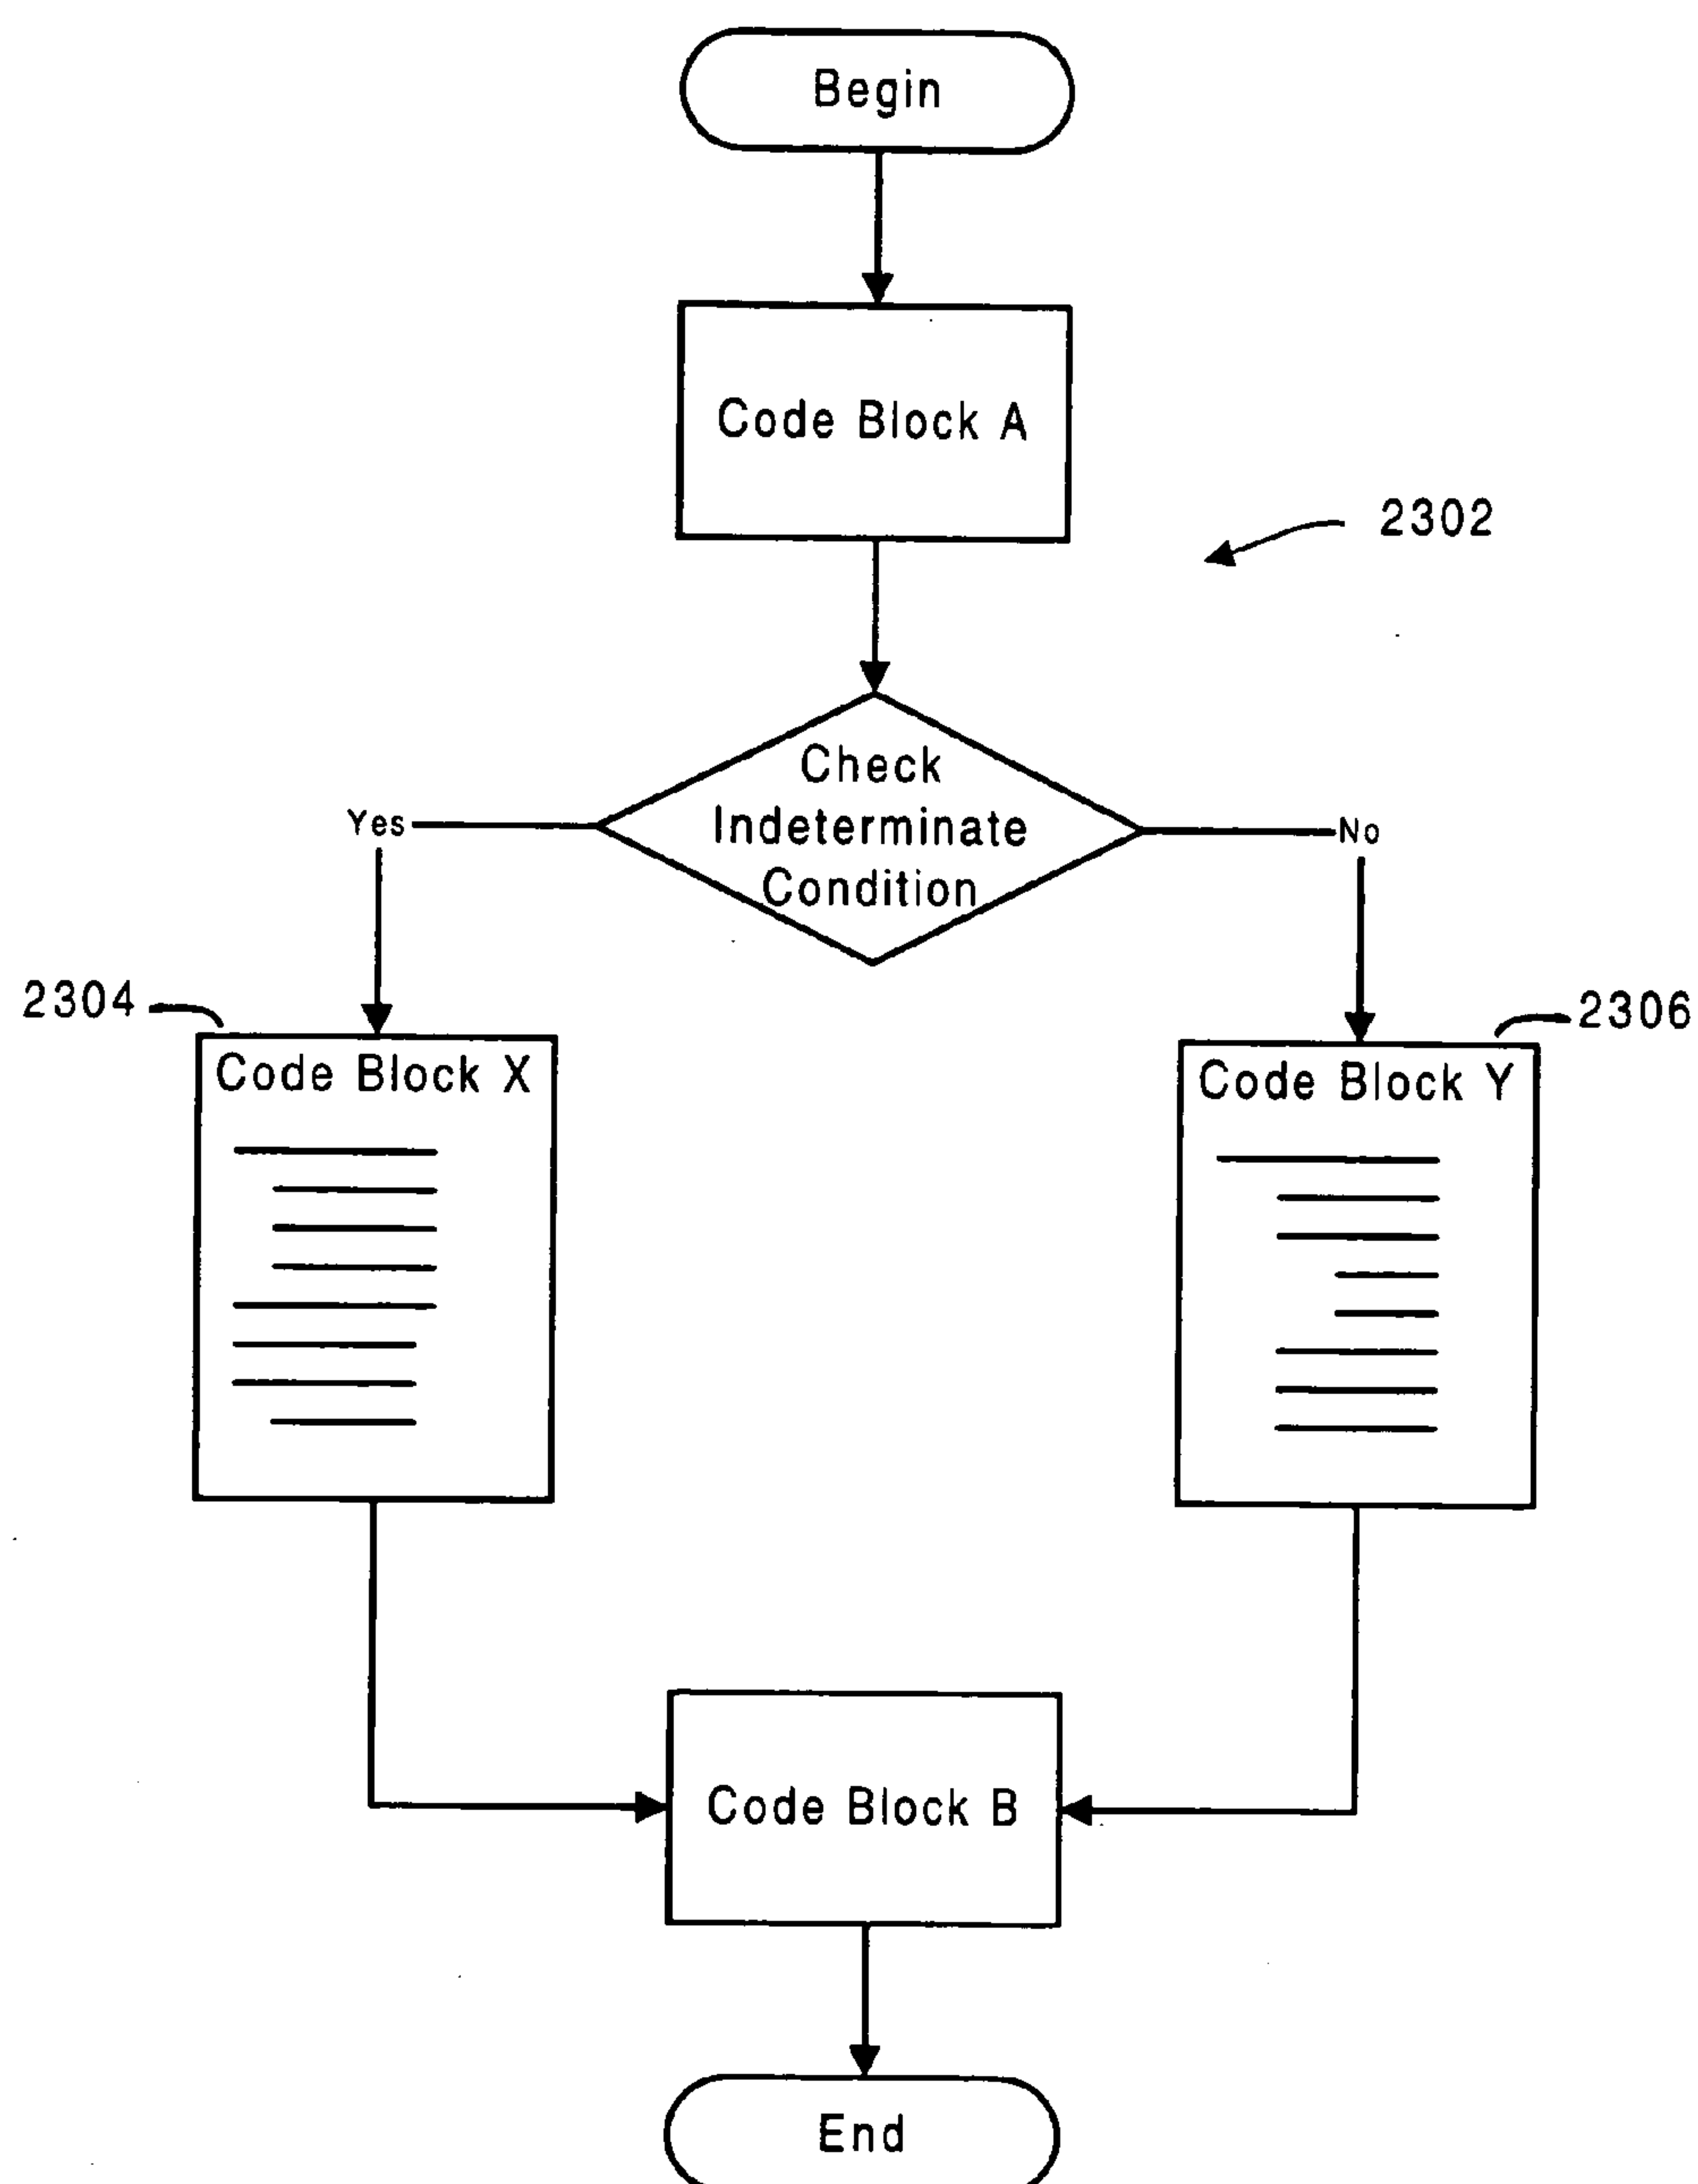


Fig. 23

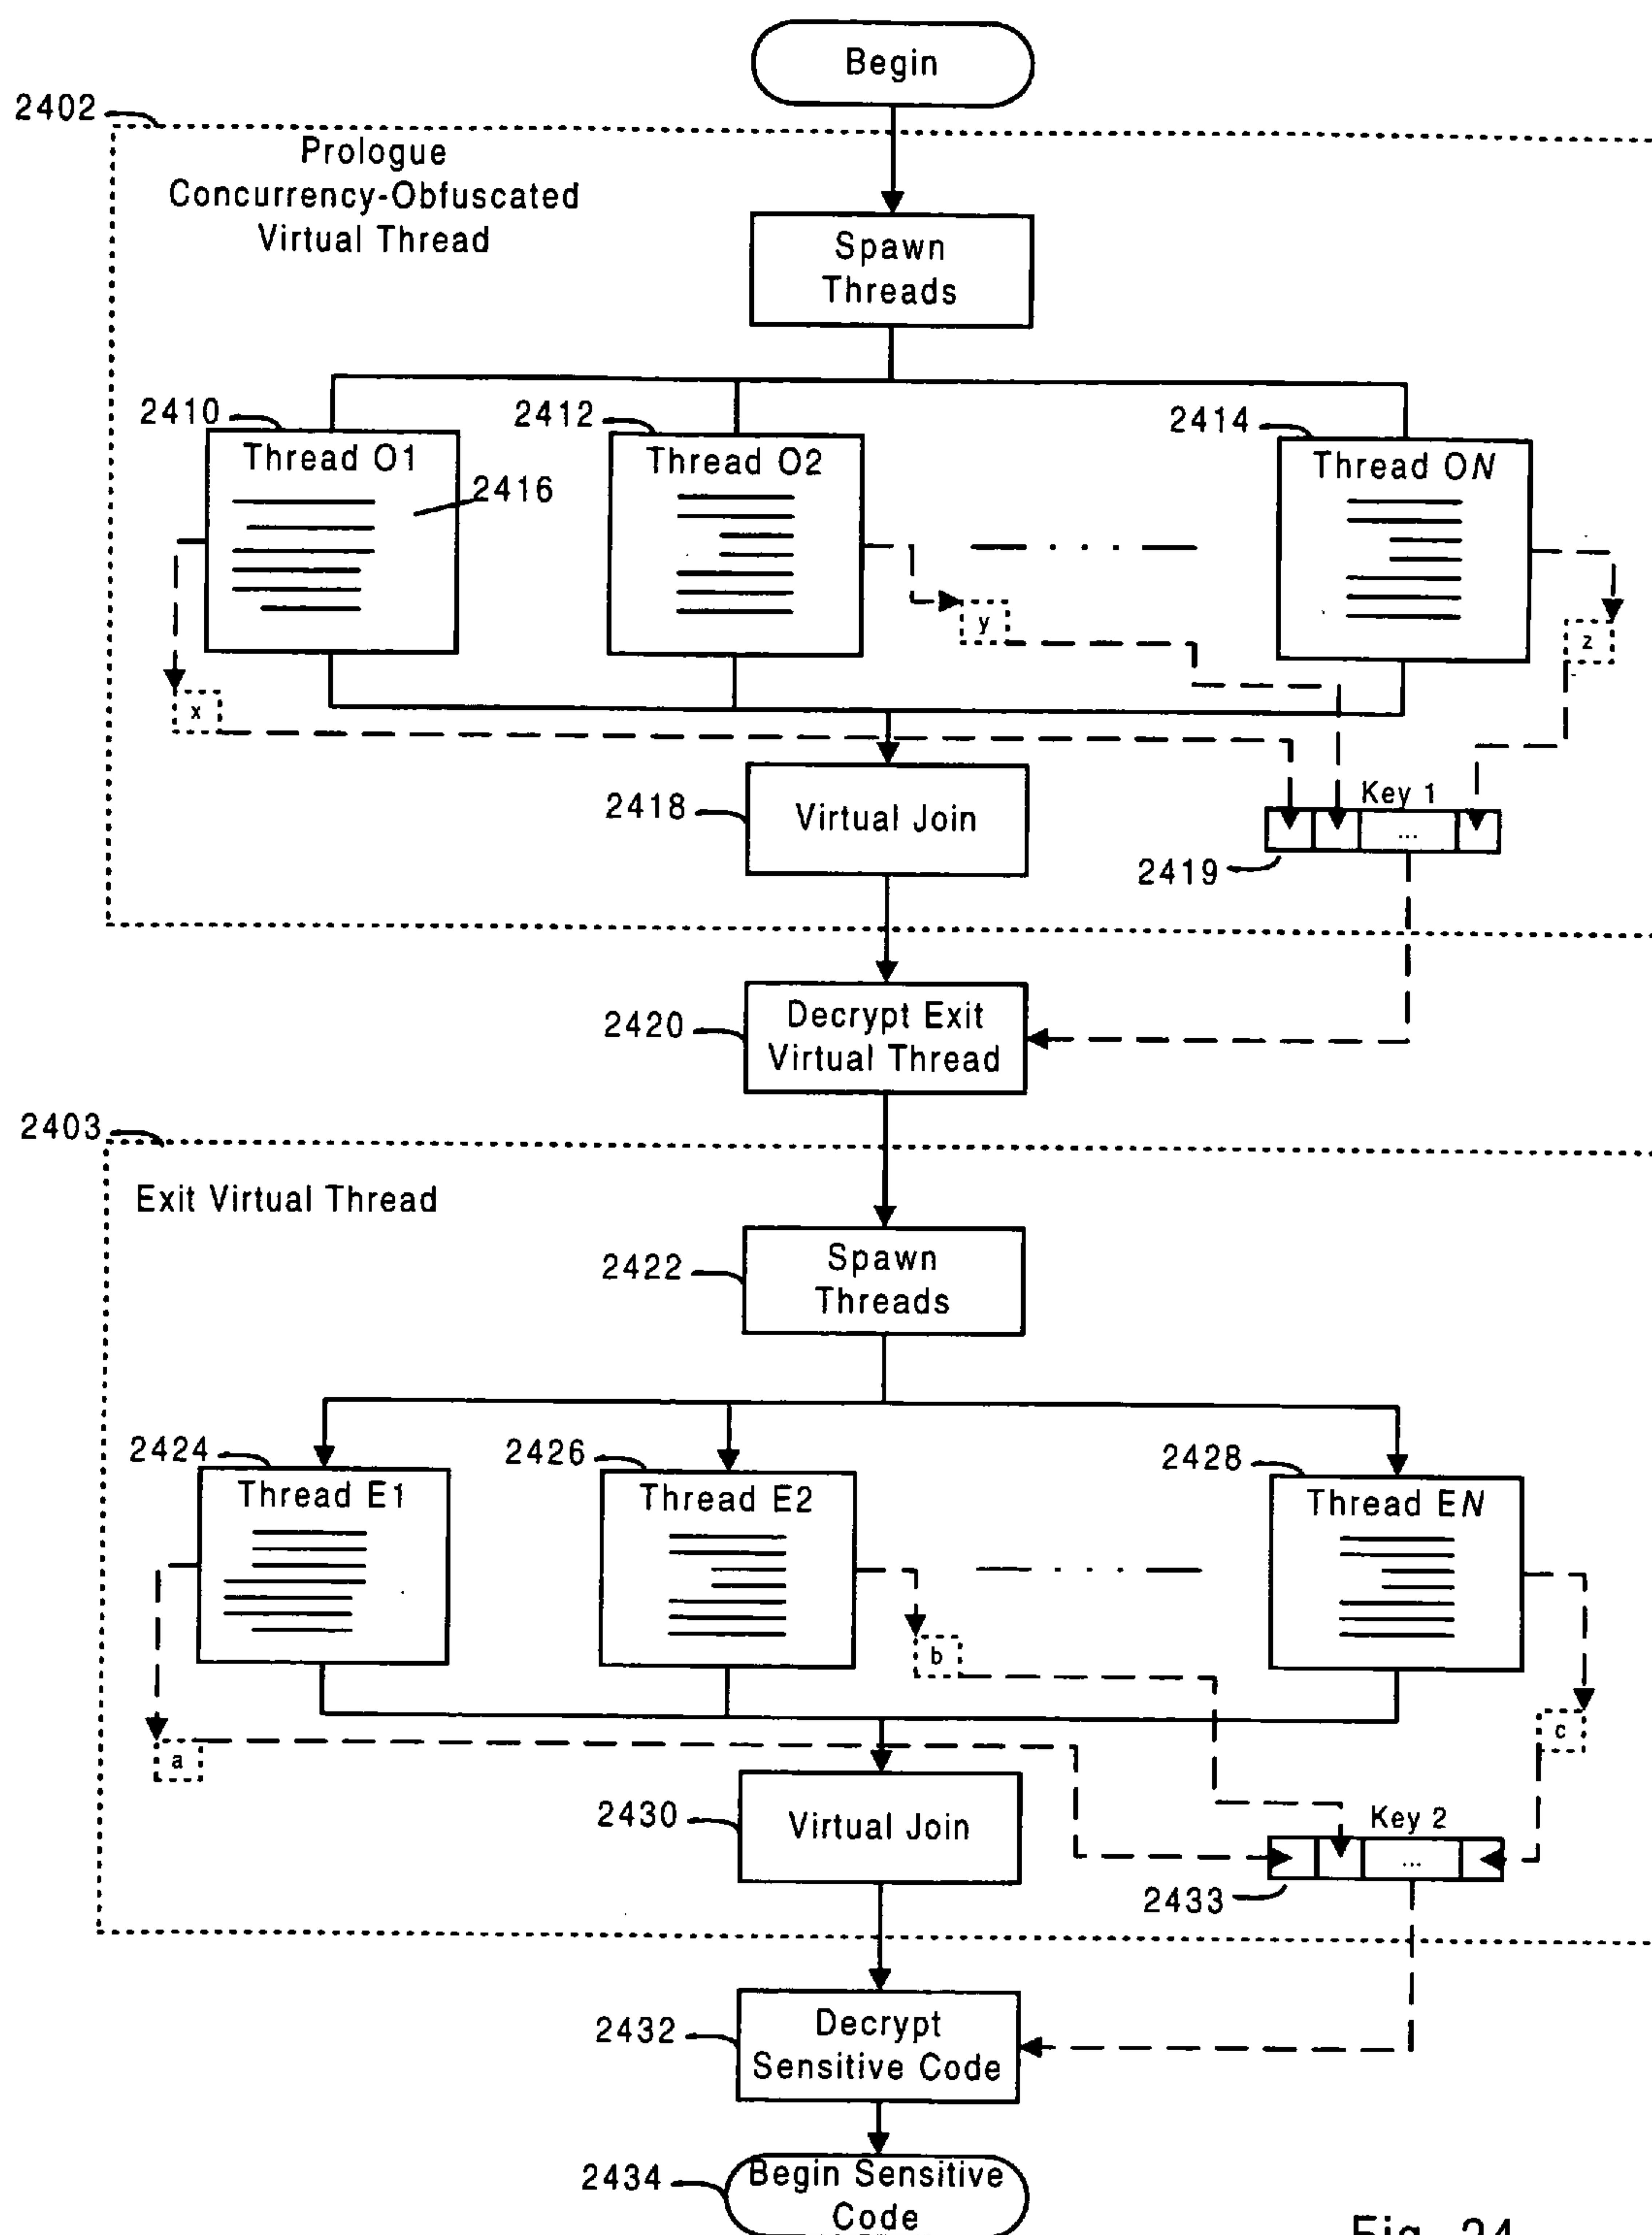


Fig. 24

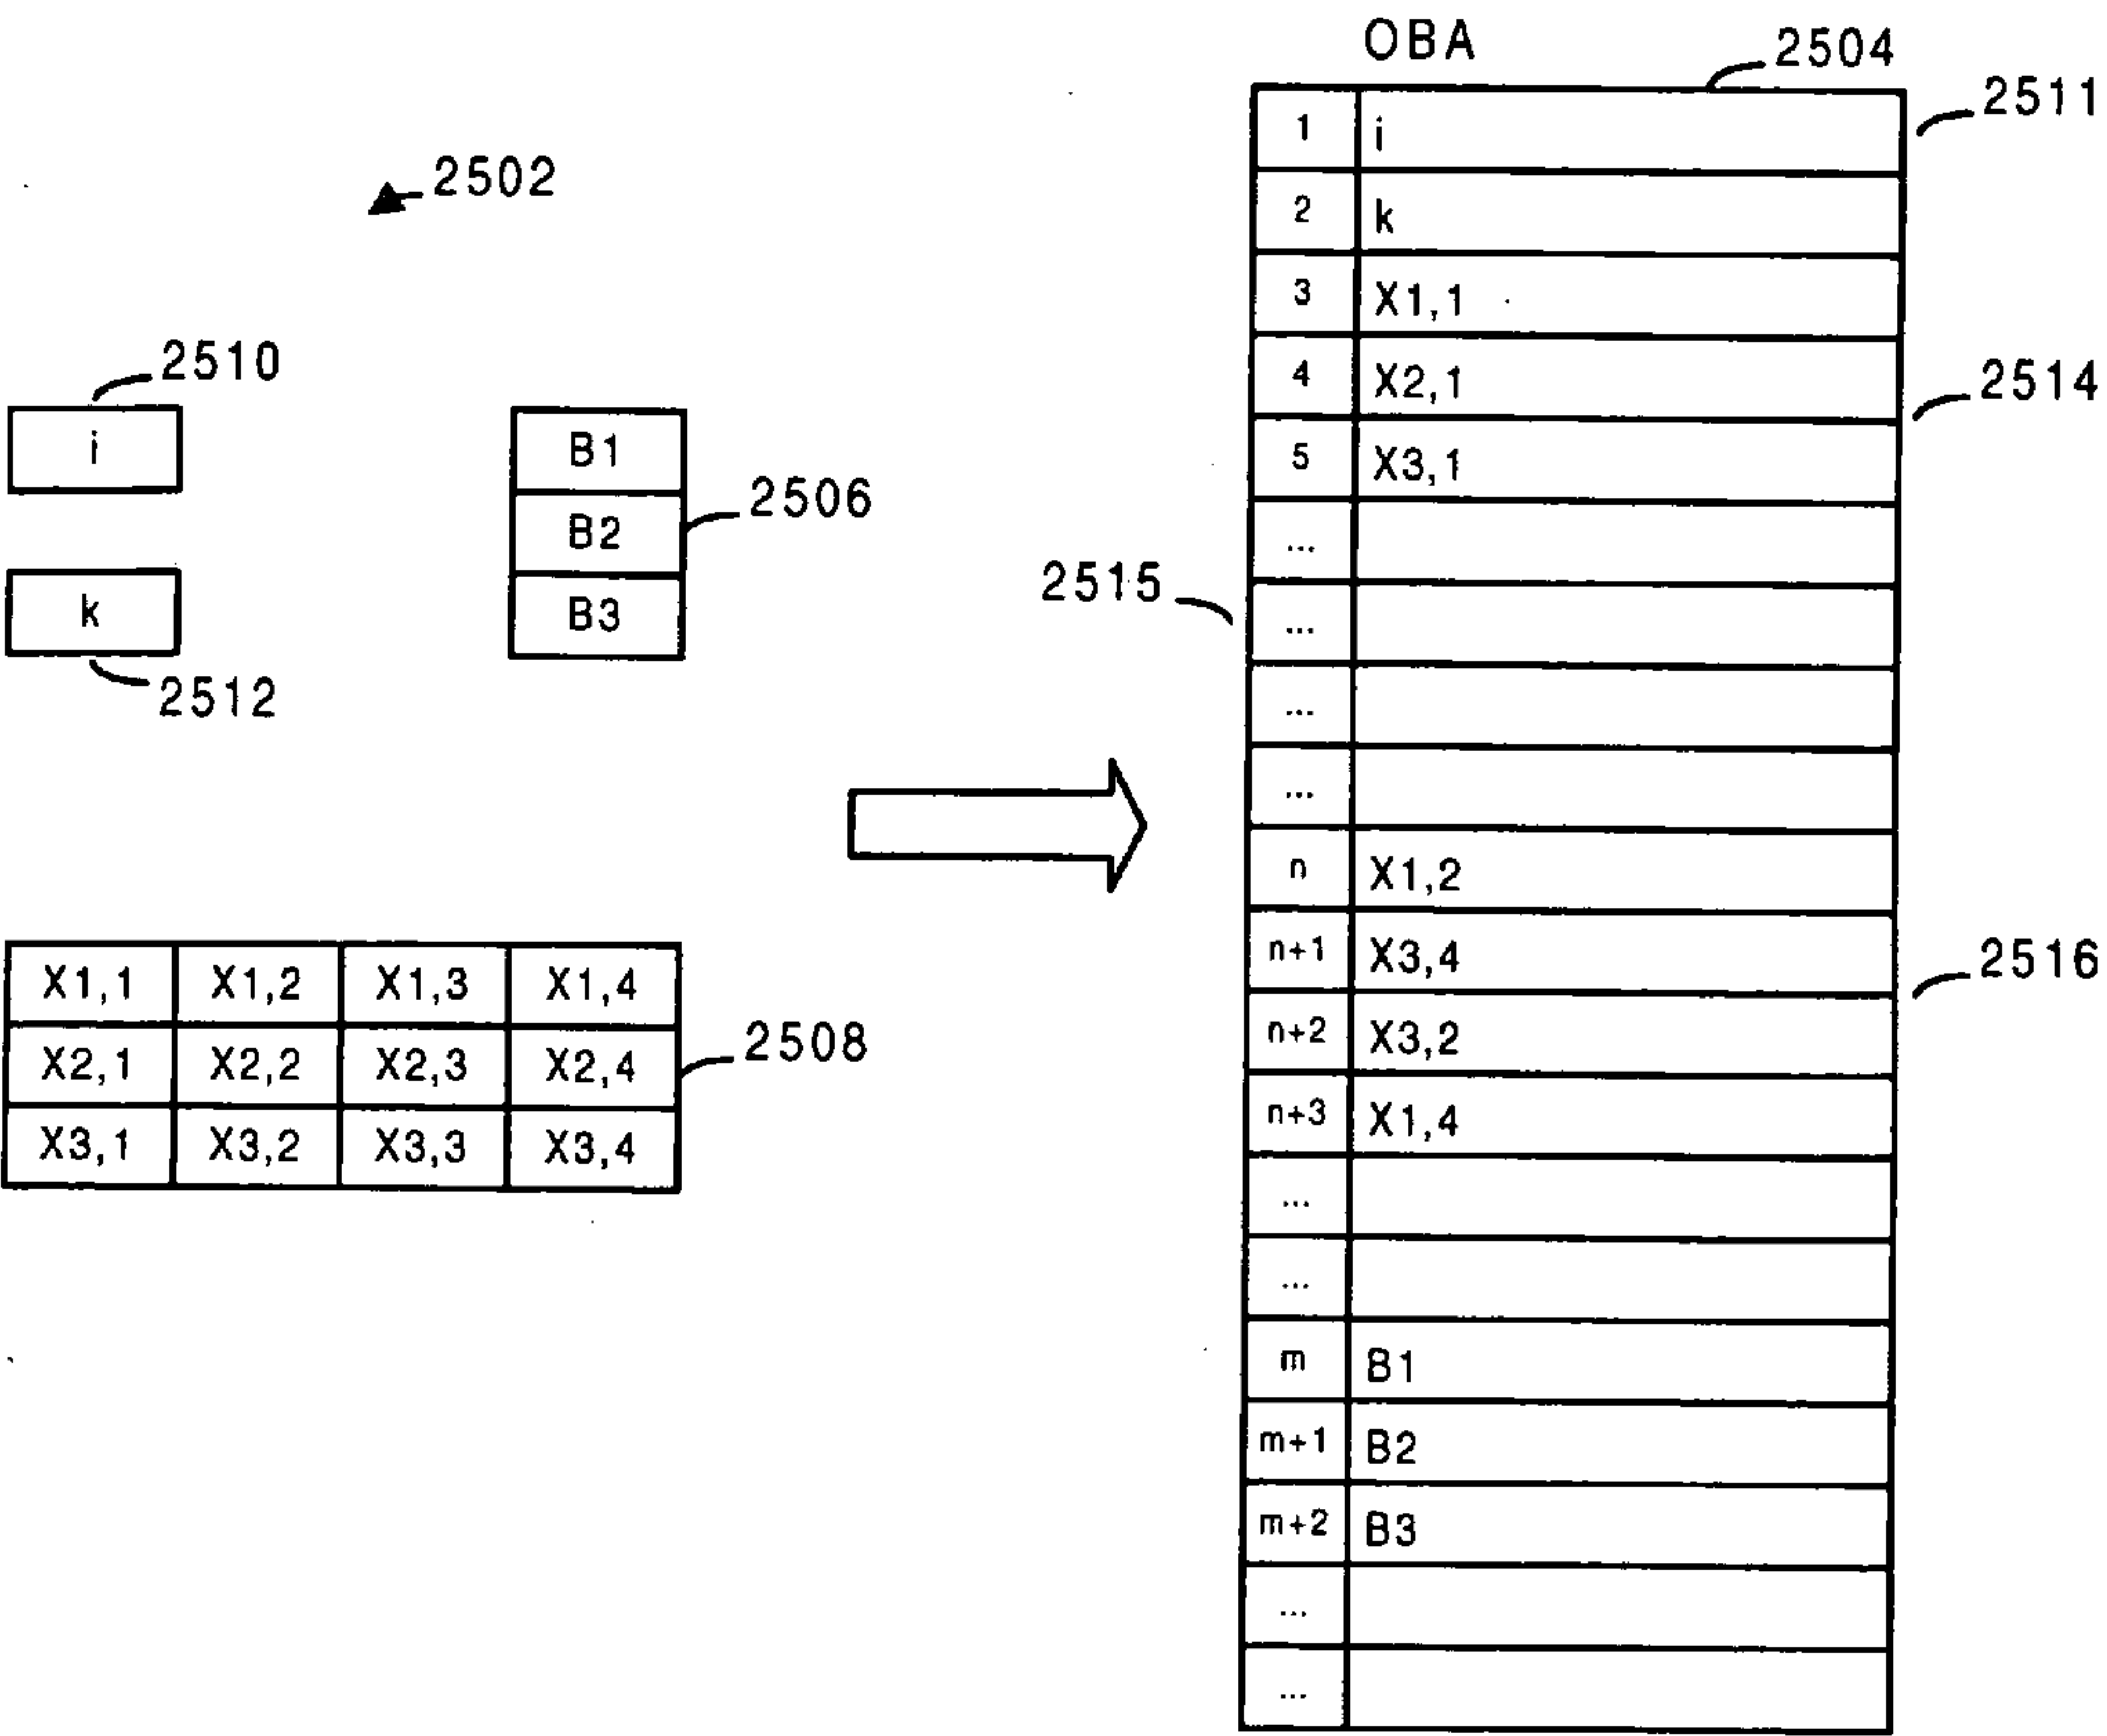


Fig. 25

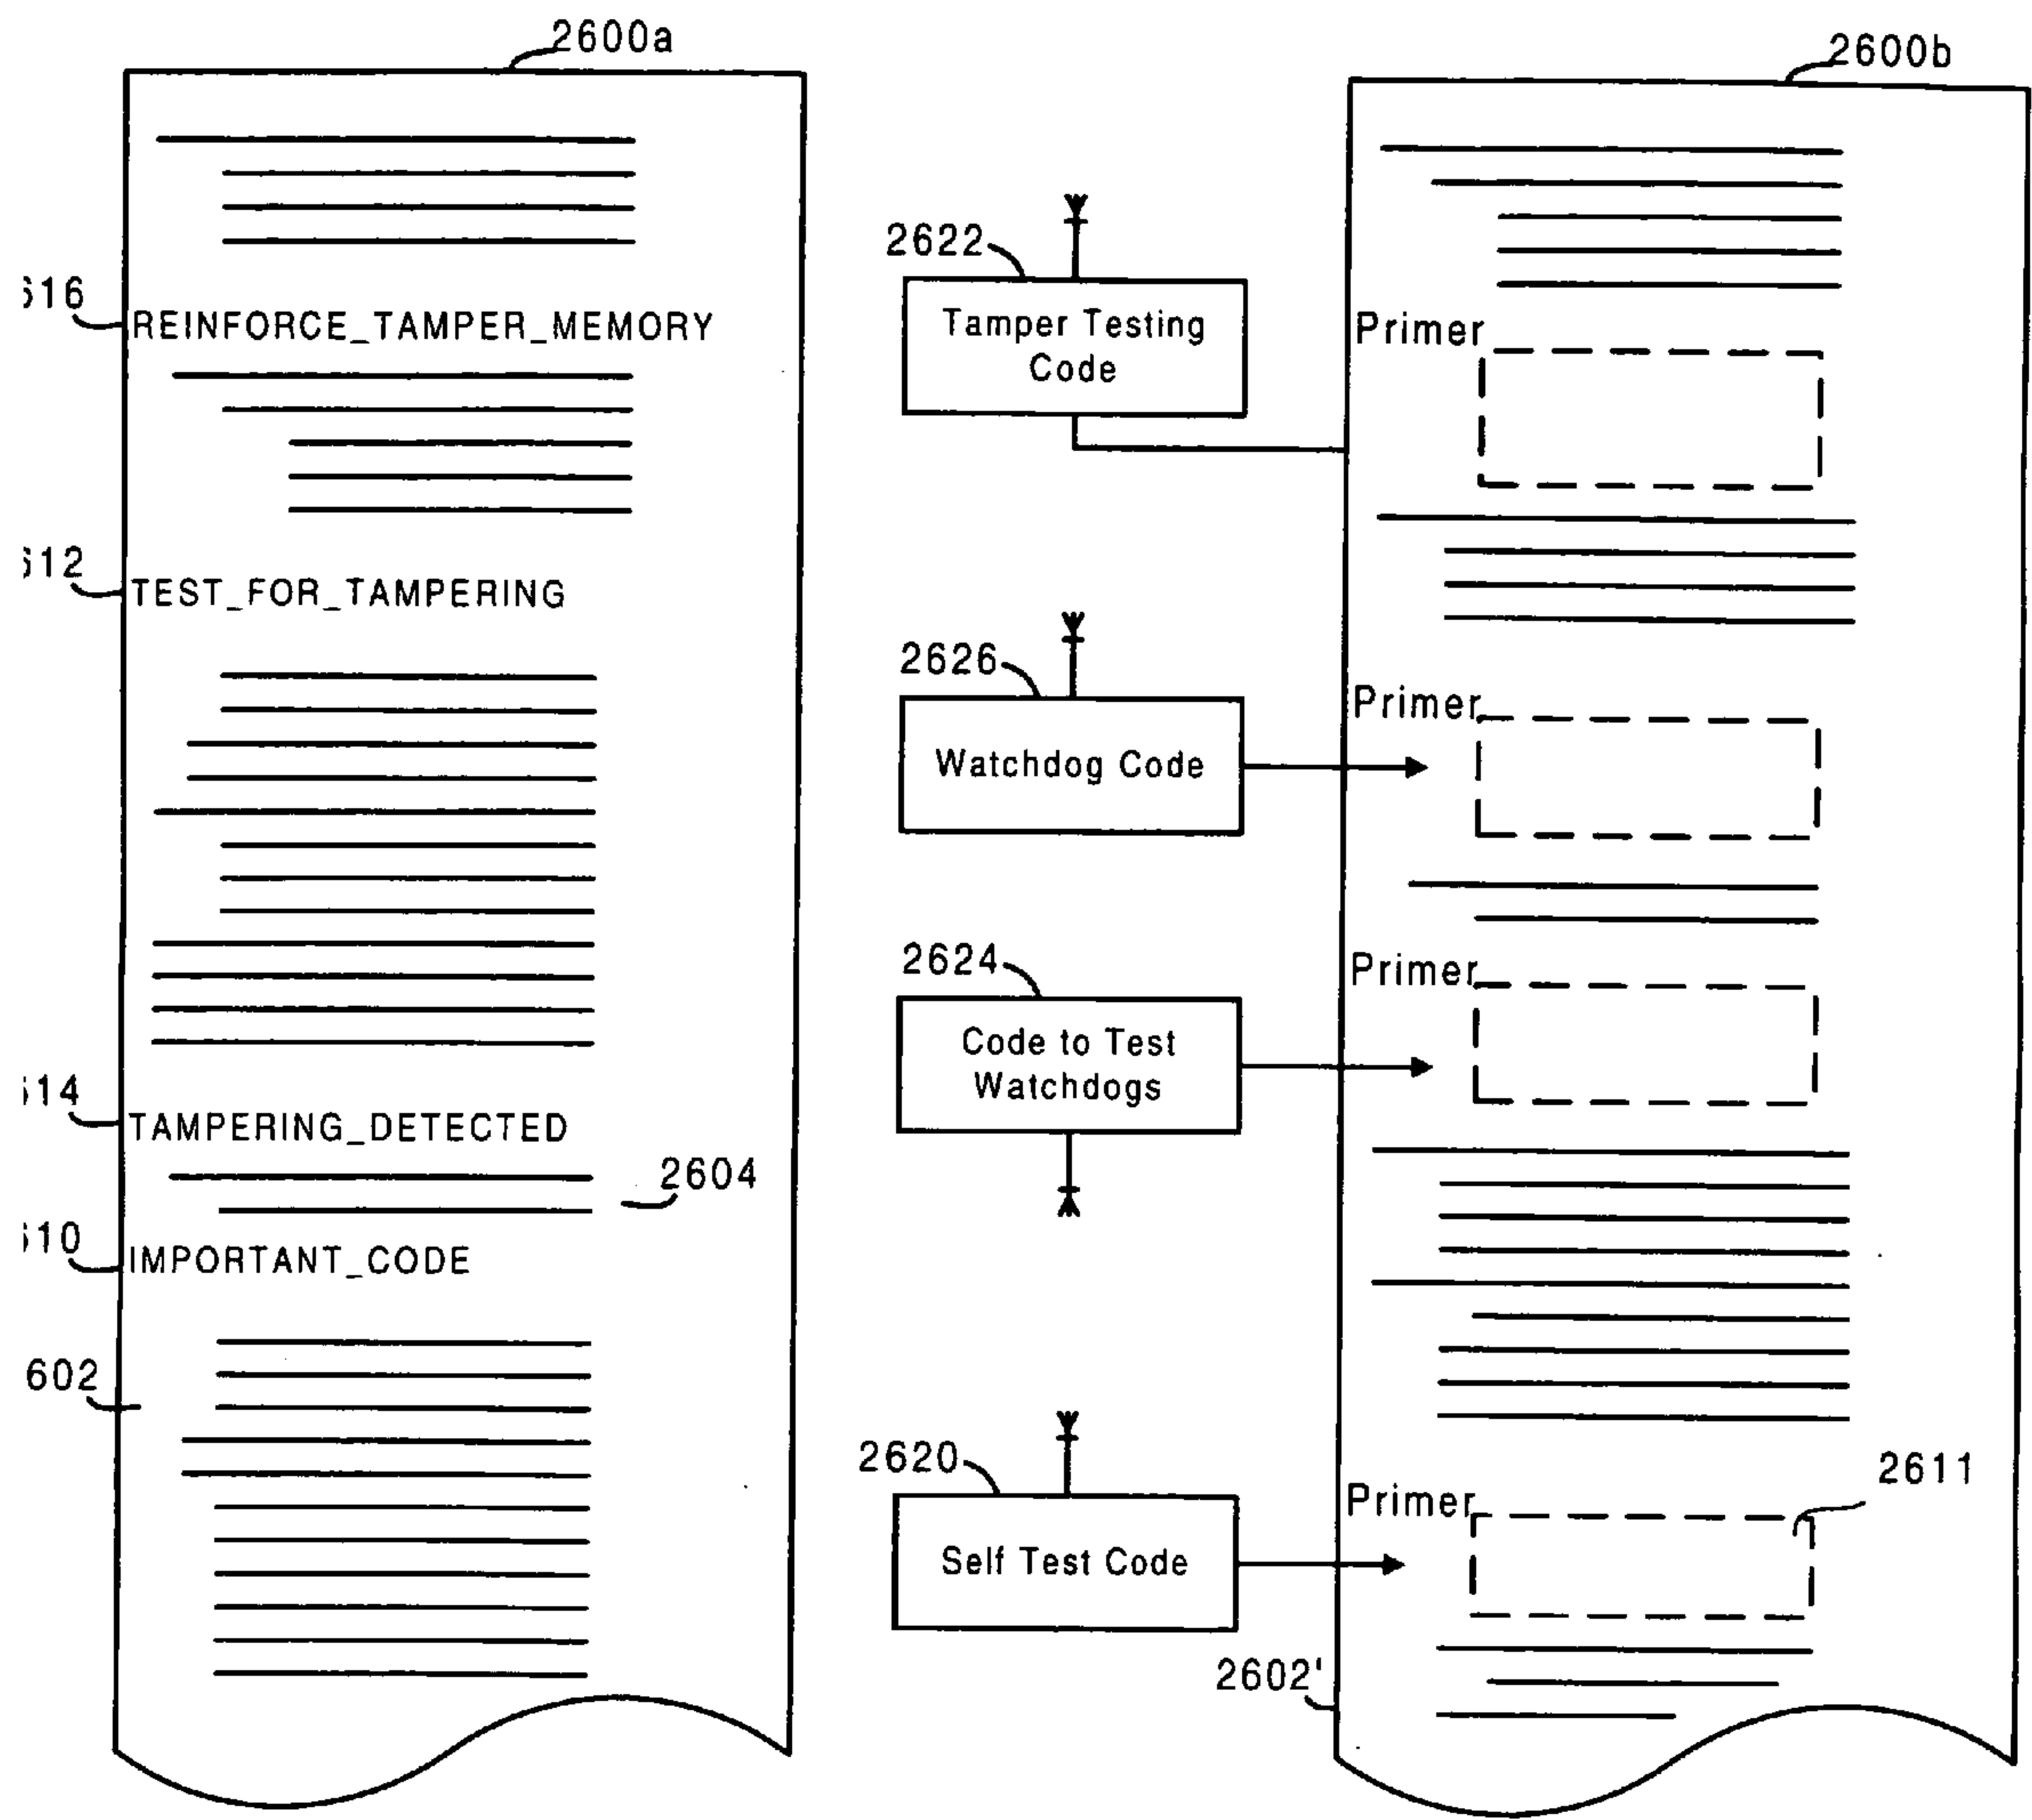


Fig. 26




Fig. 27

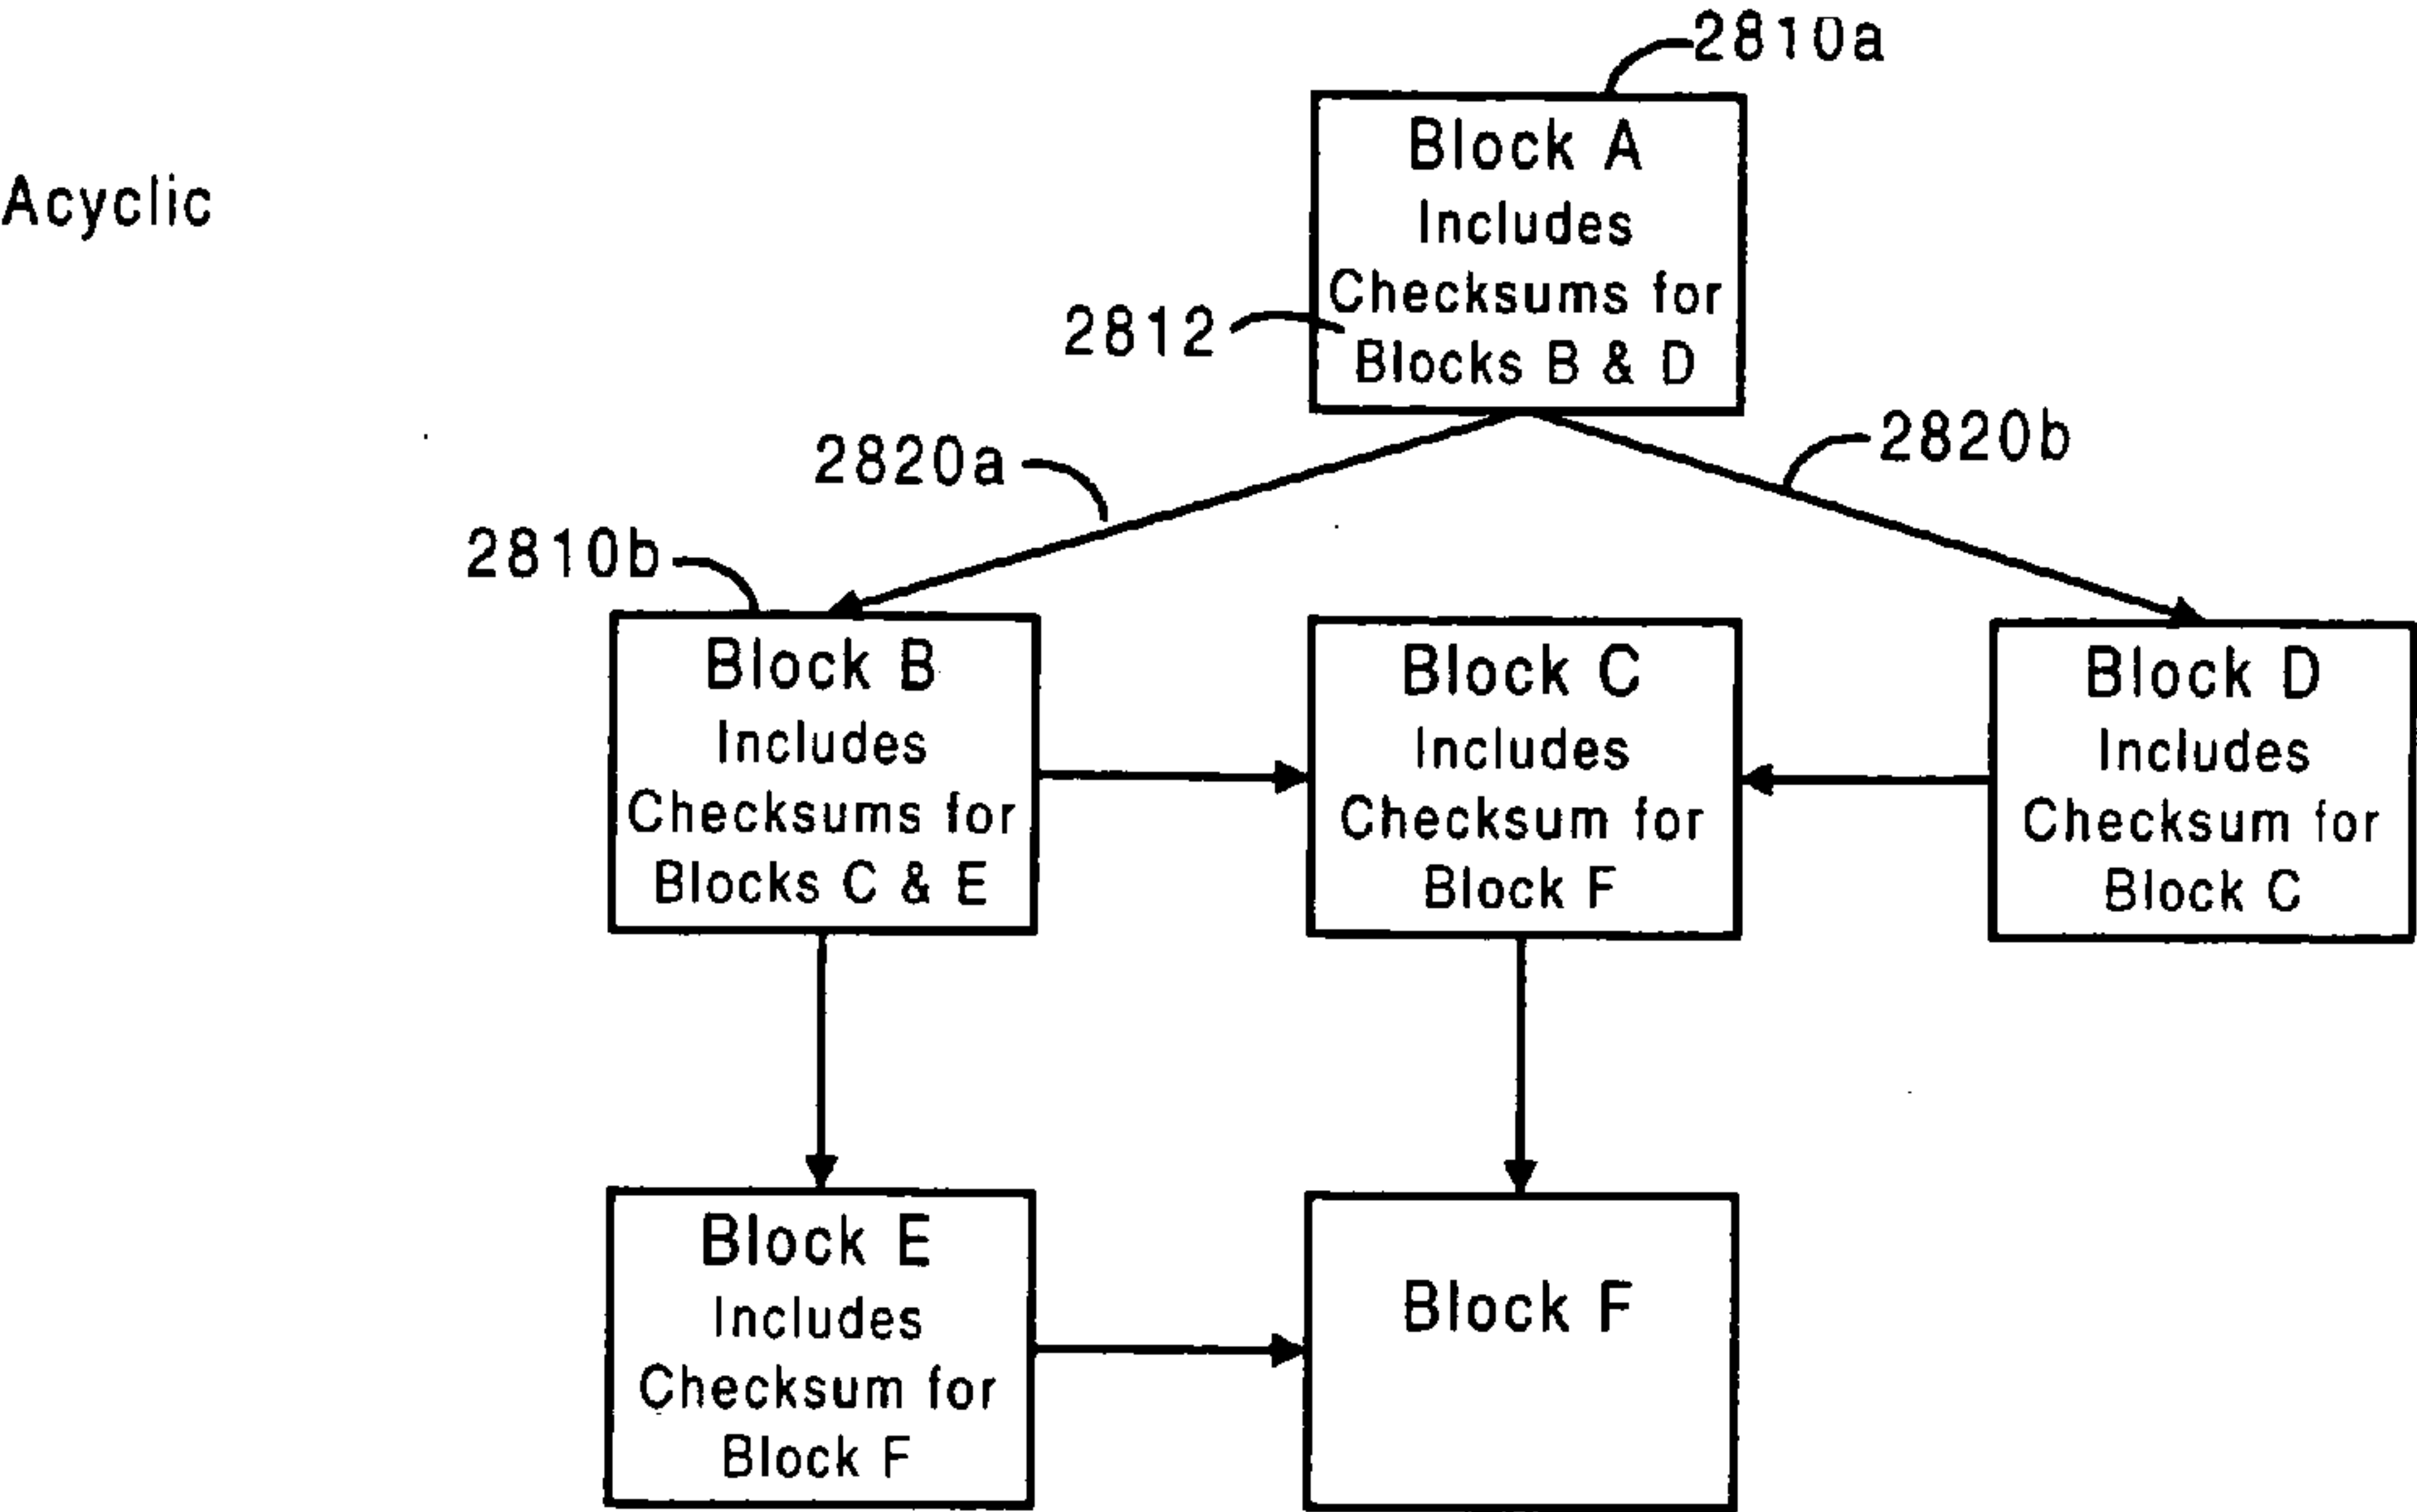


Fig. 28A

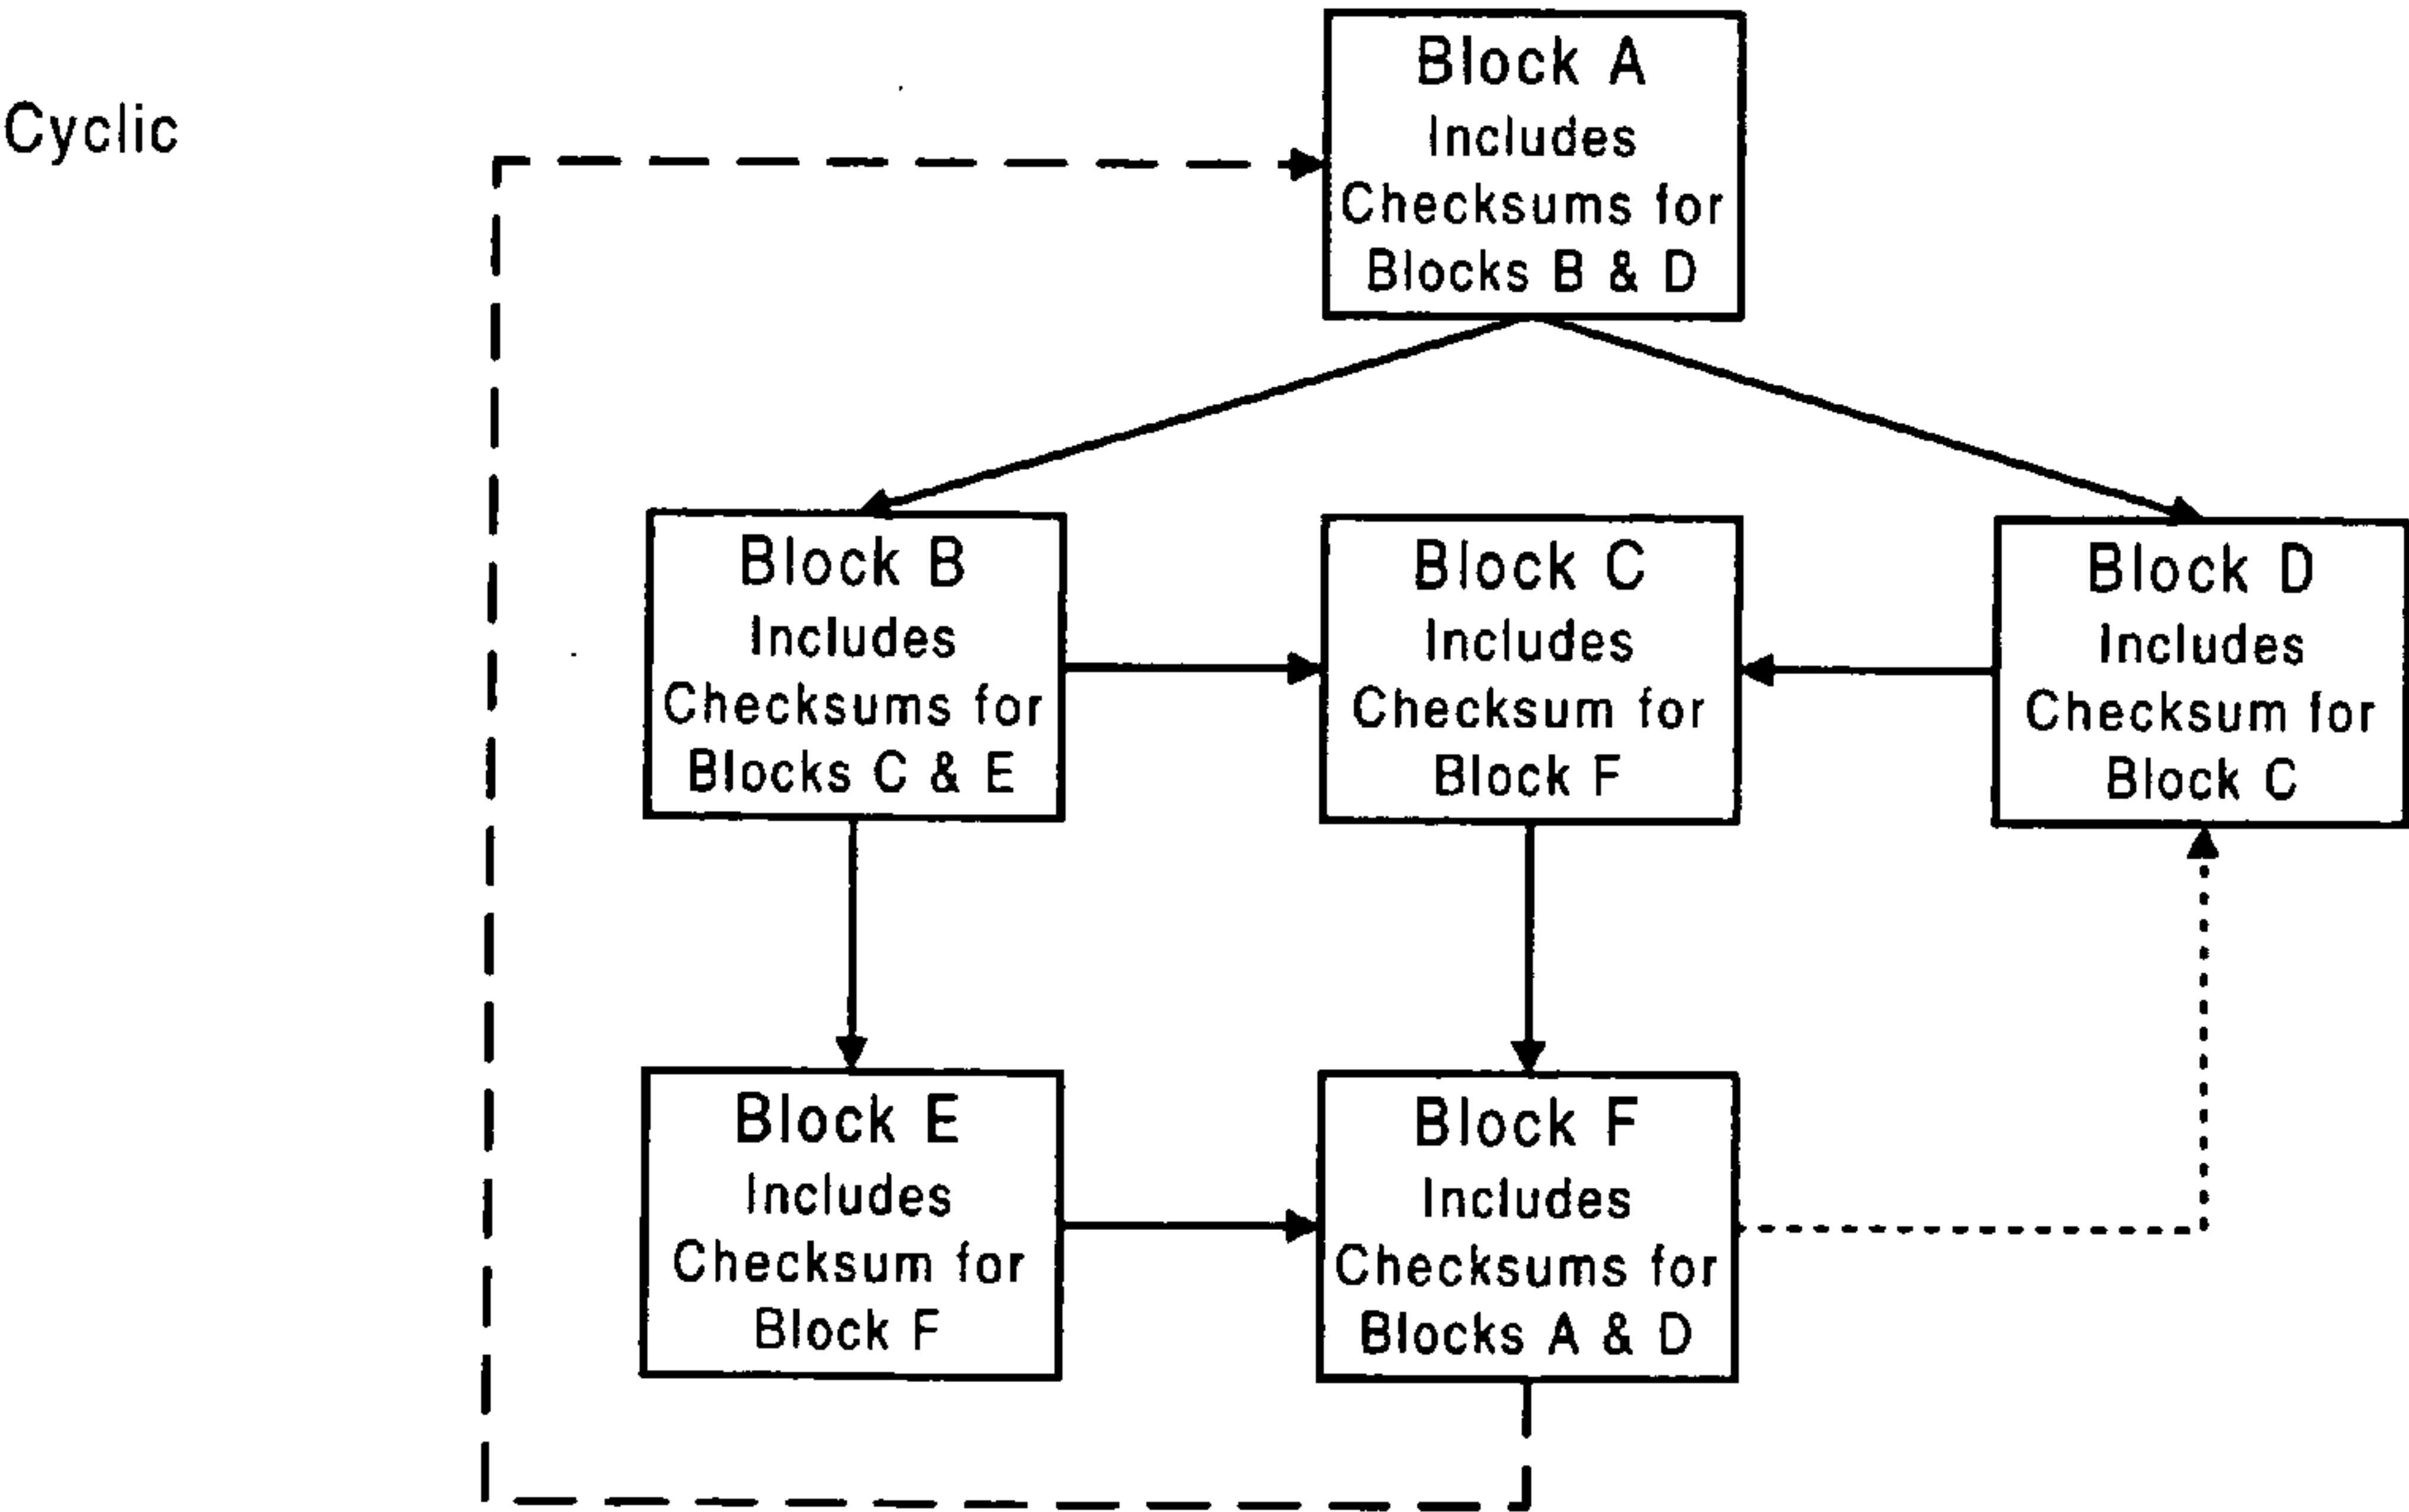


Fig. 28B

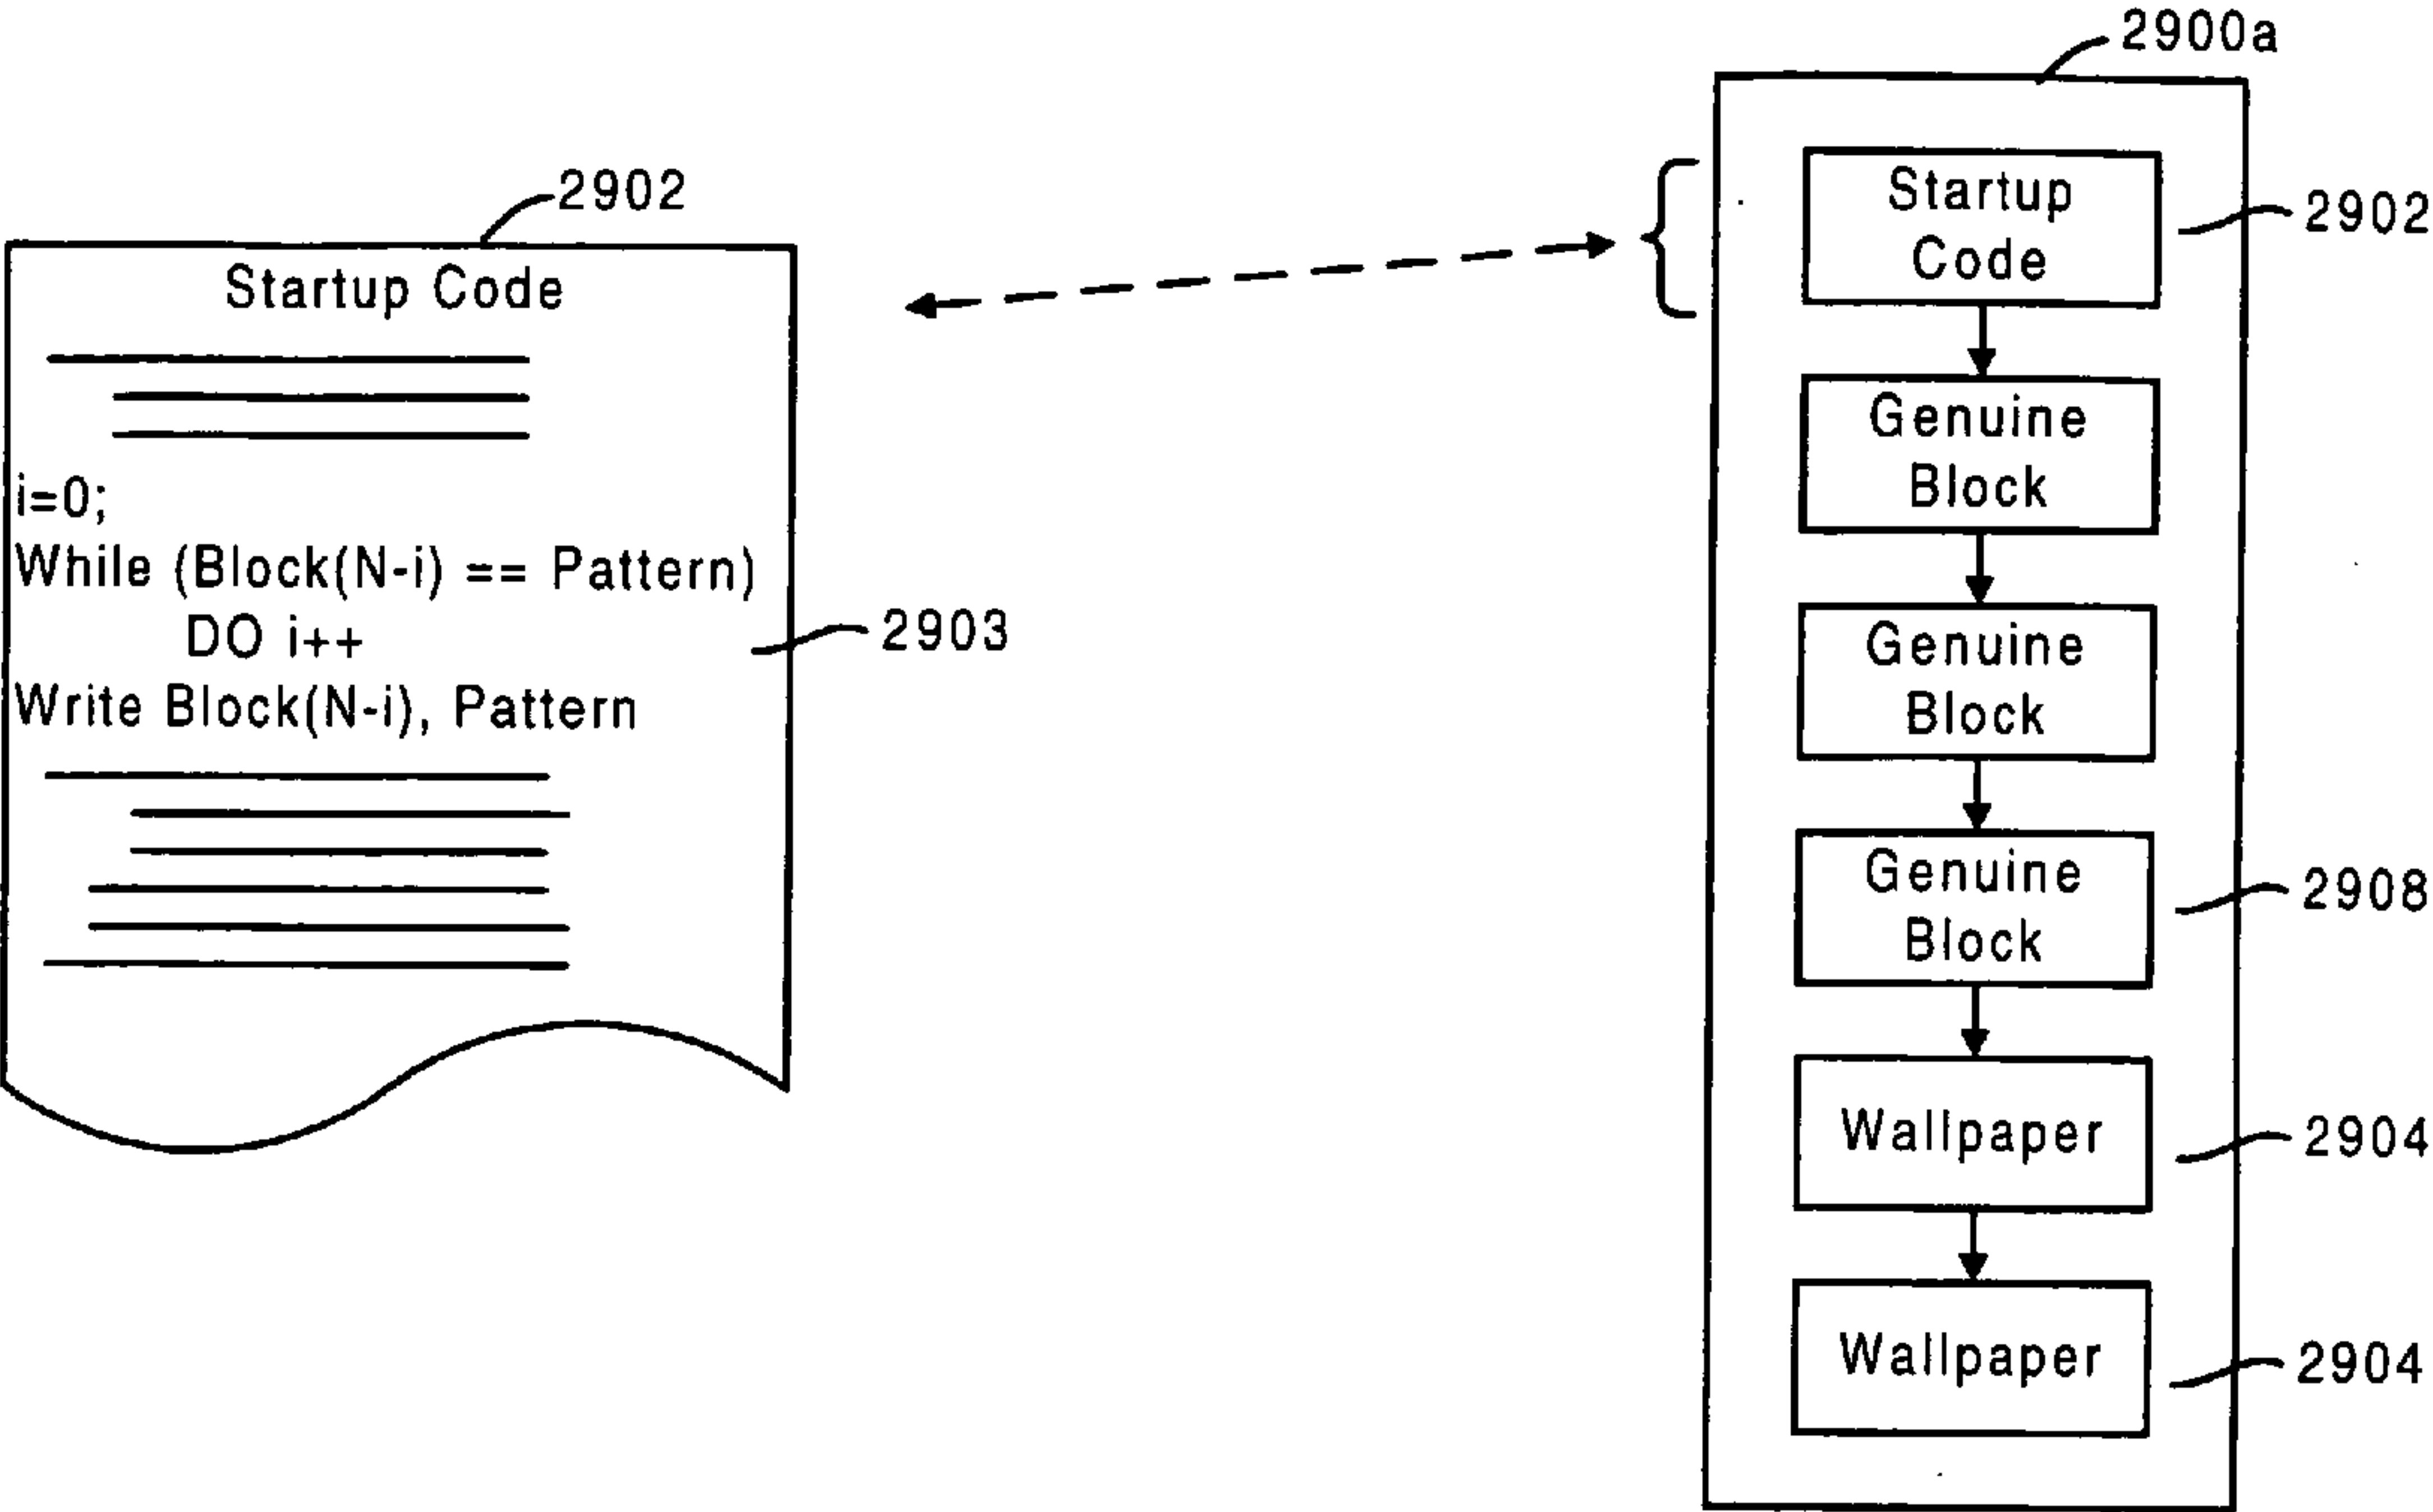


Fig. 29A

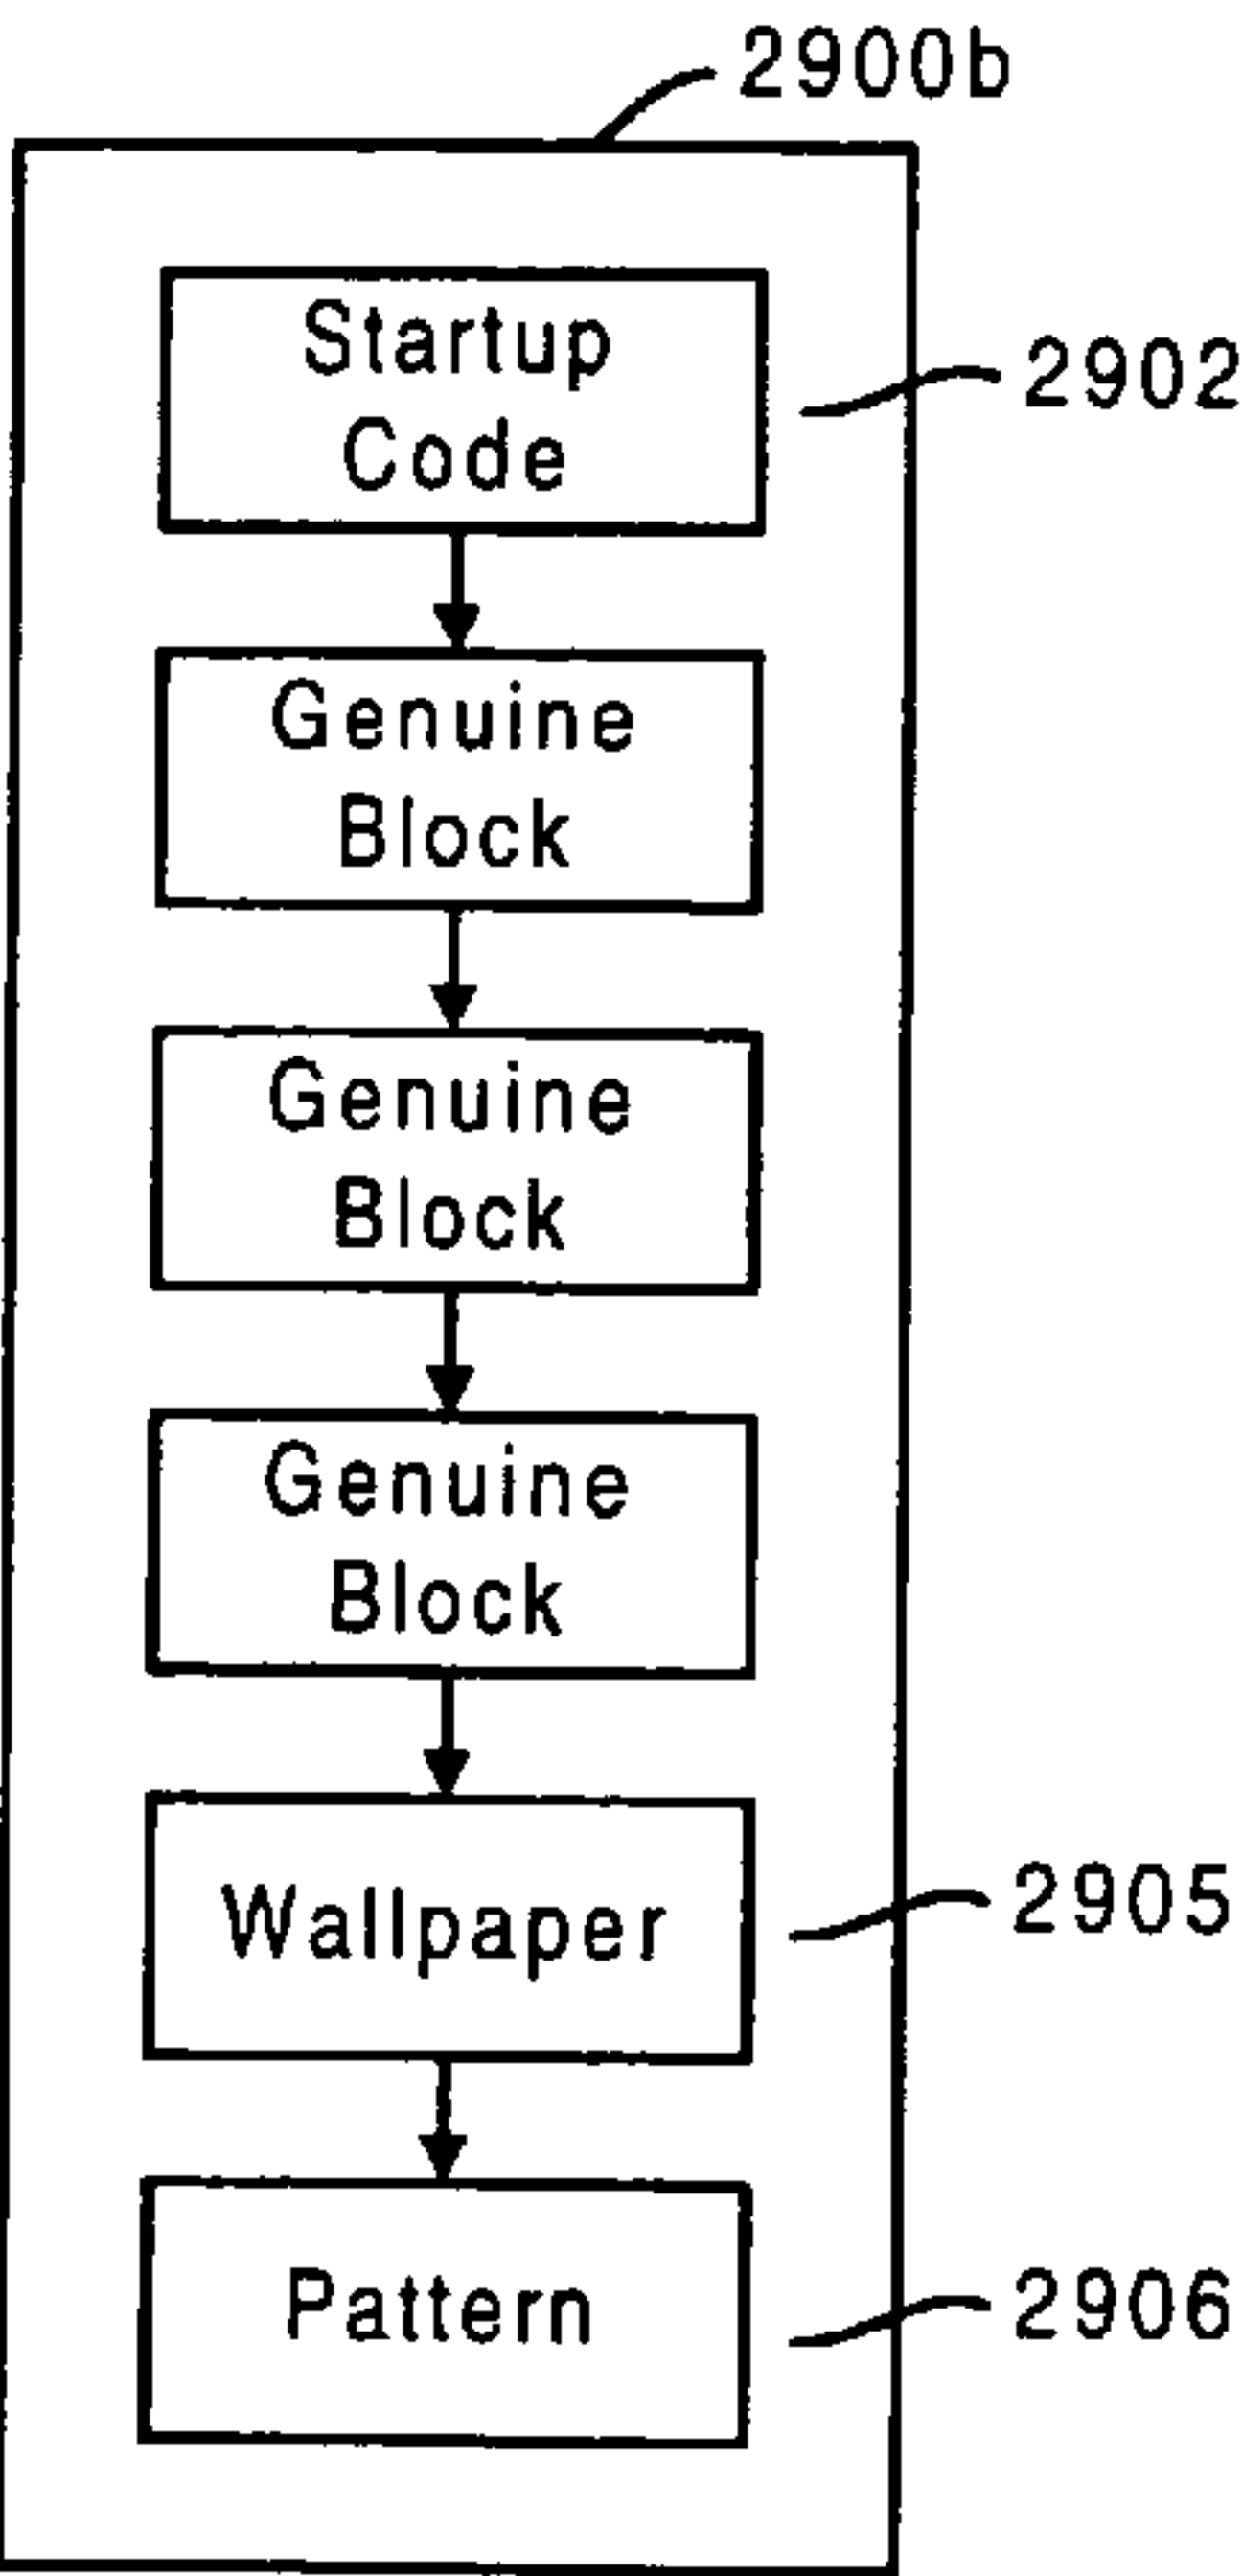


Fig. 29B




Fig. 29C

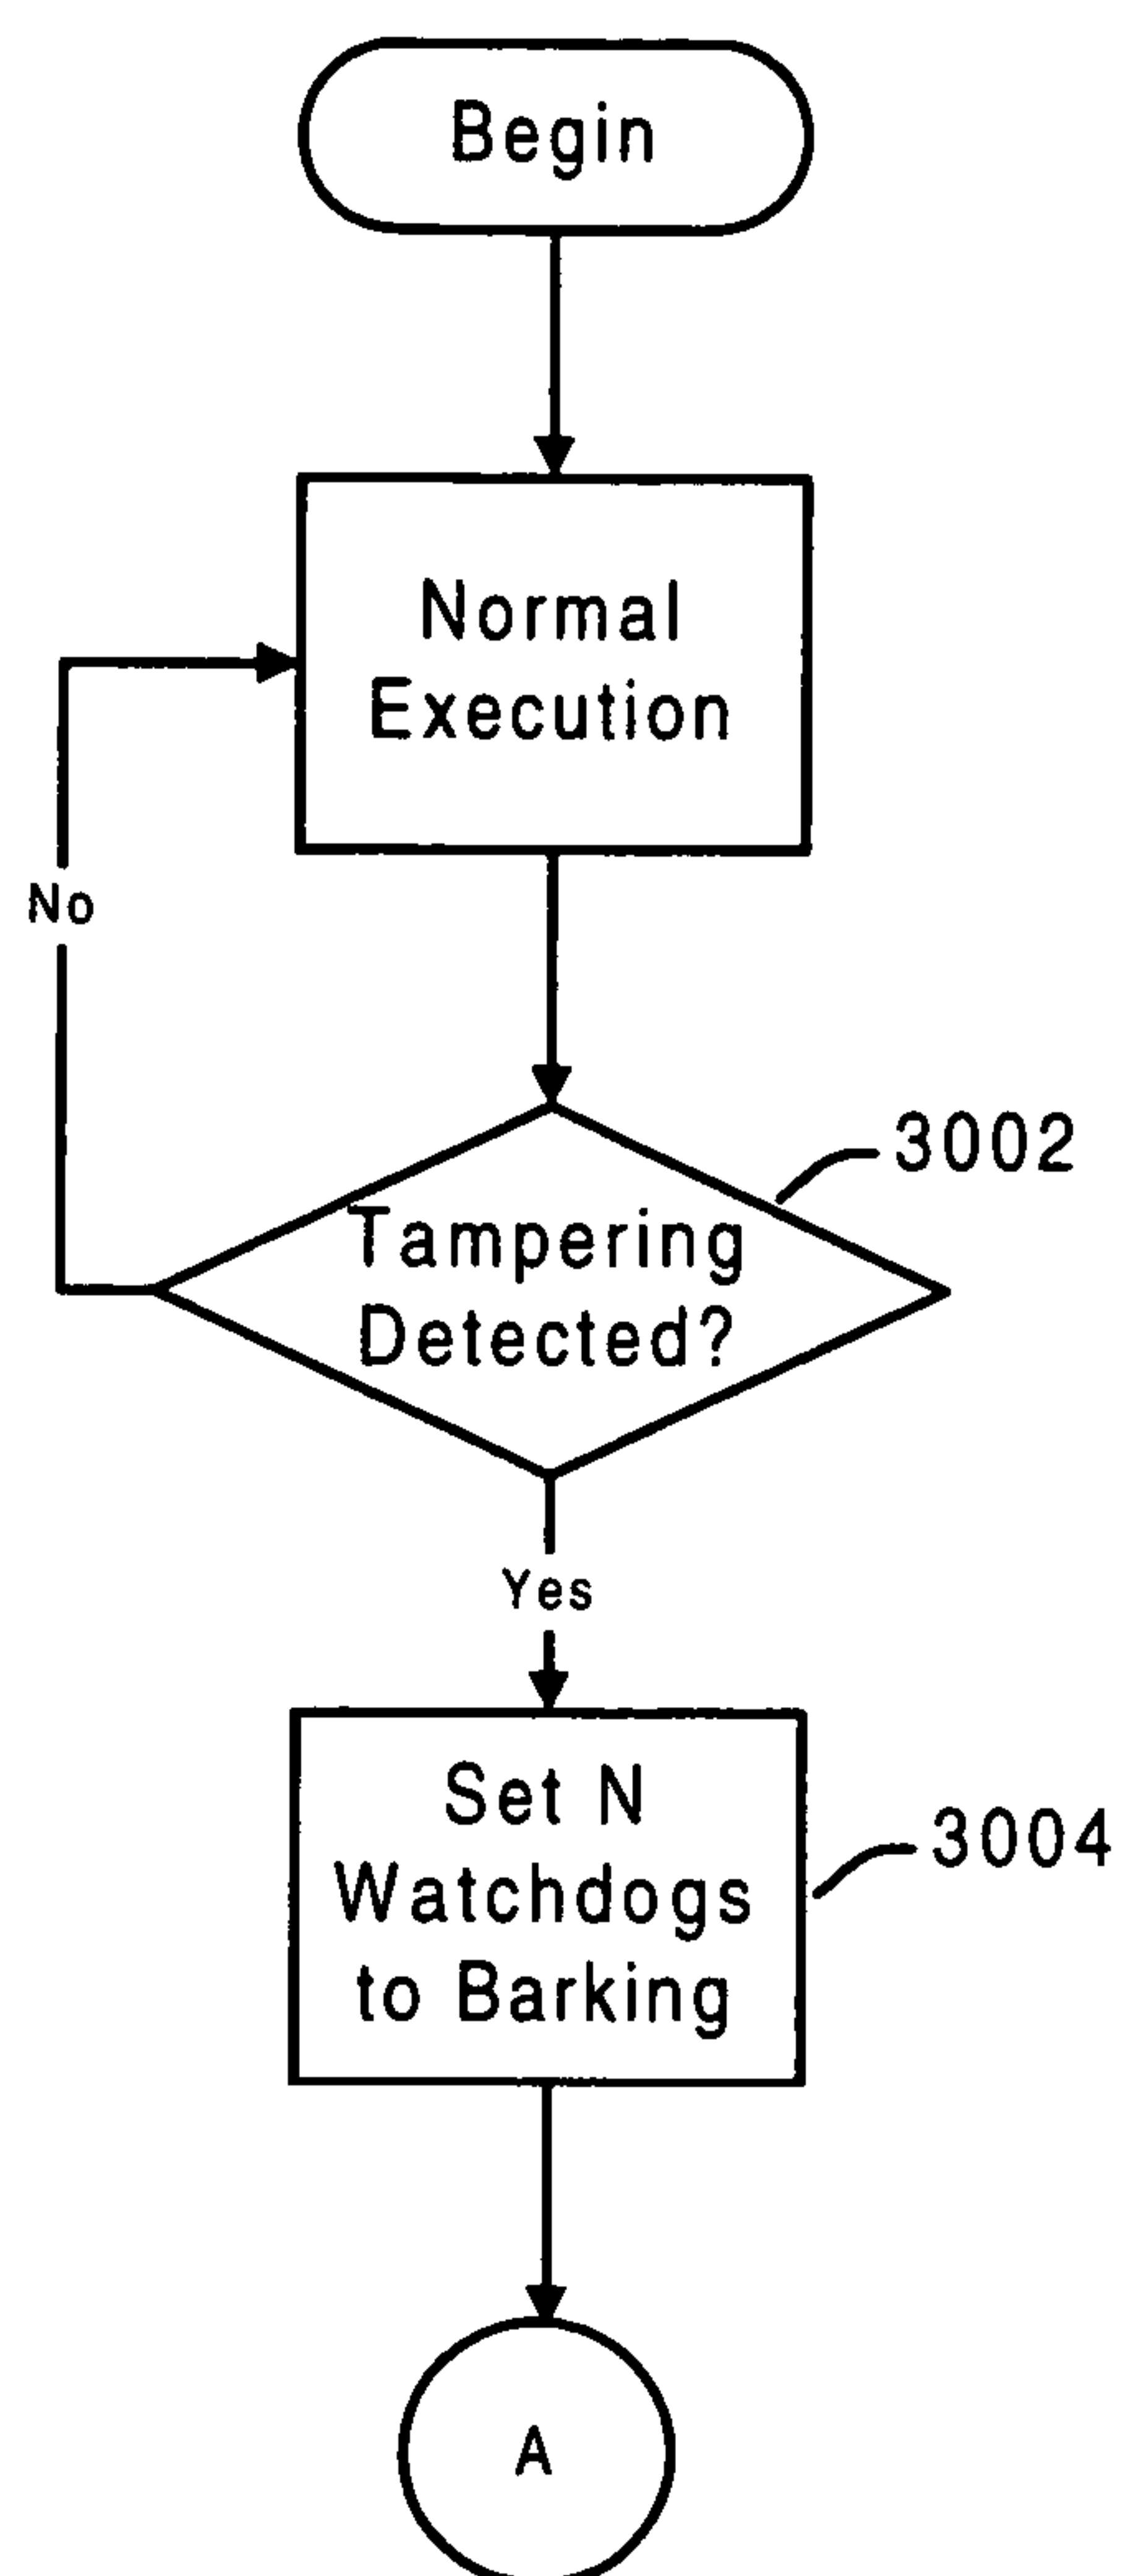


Fig. 30A

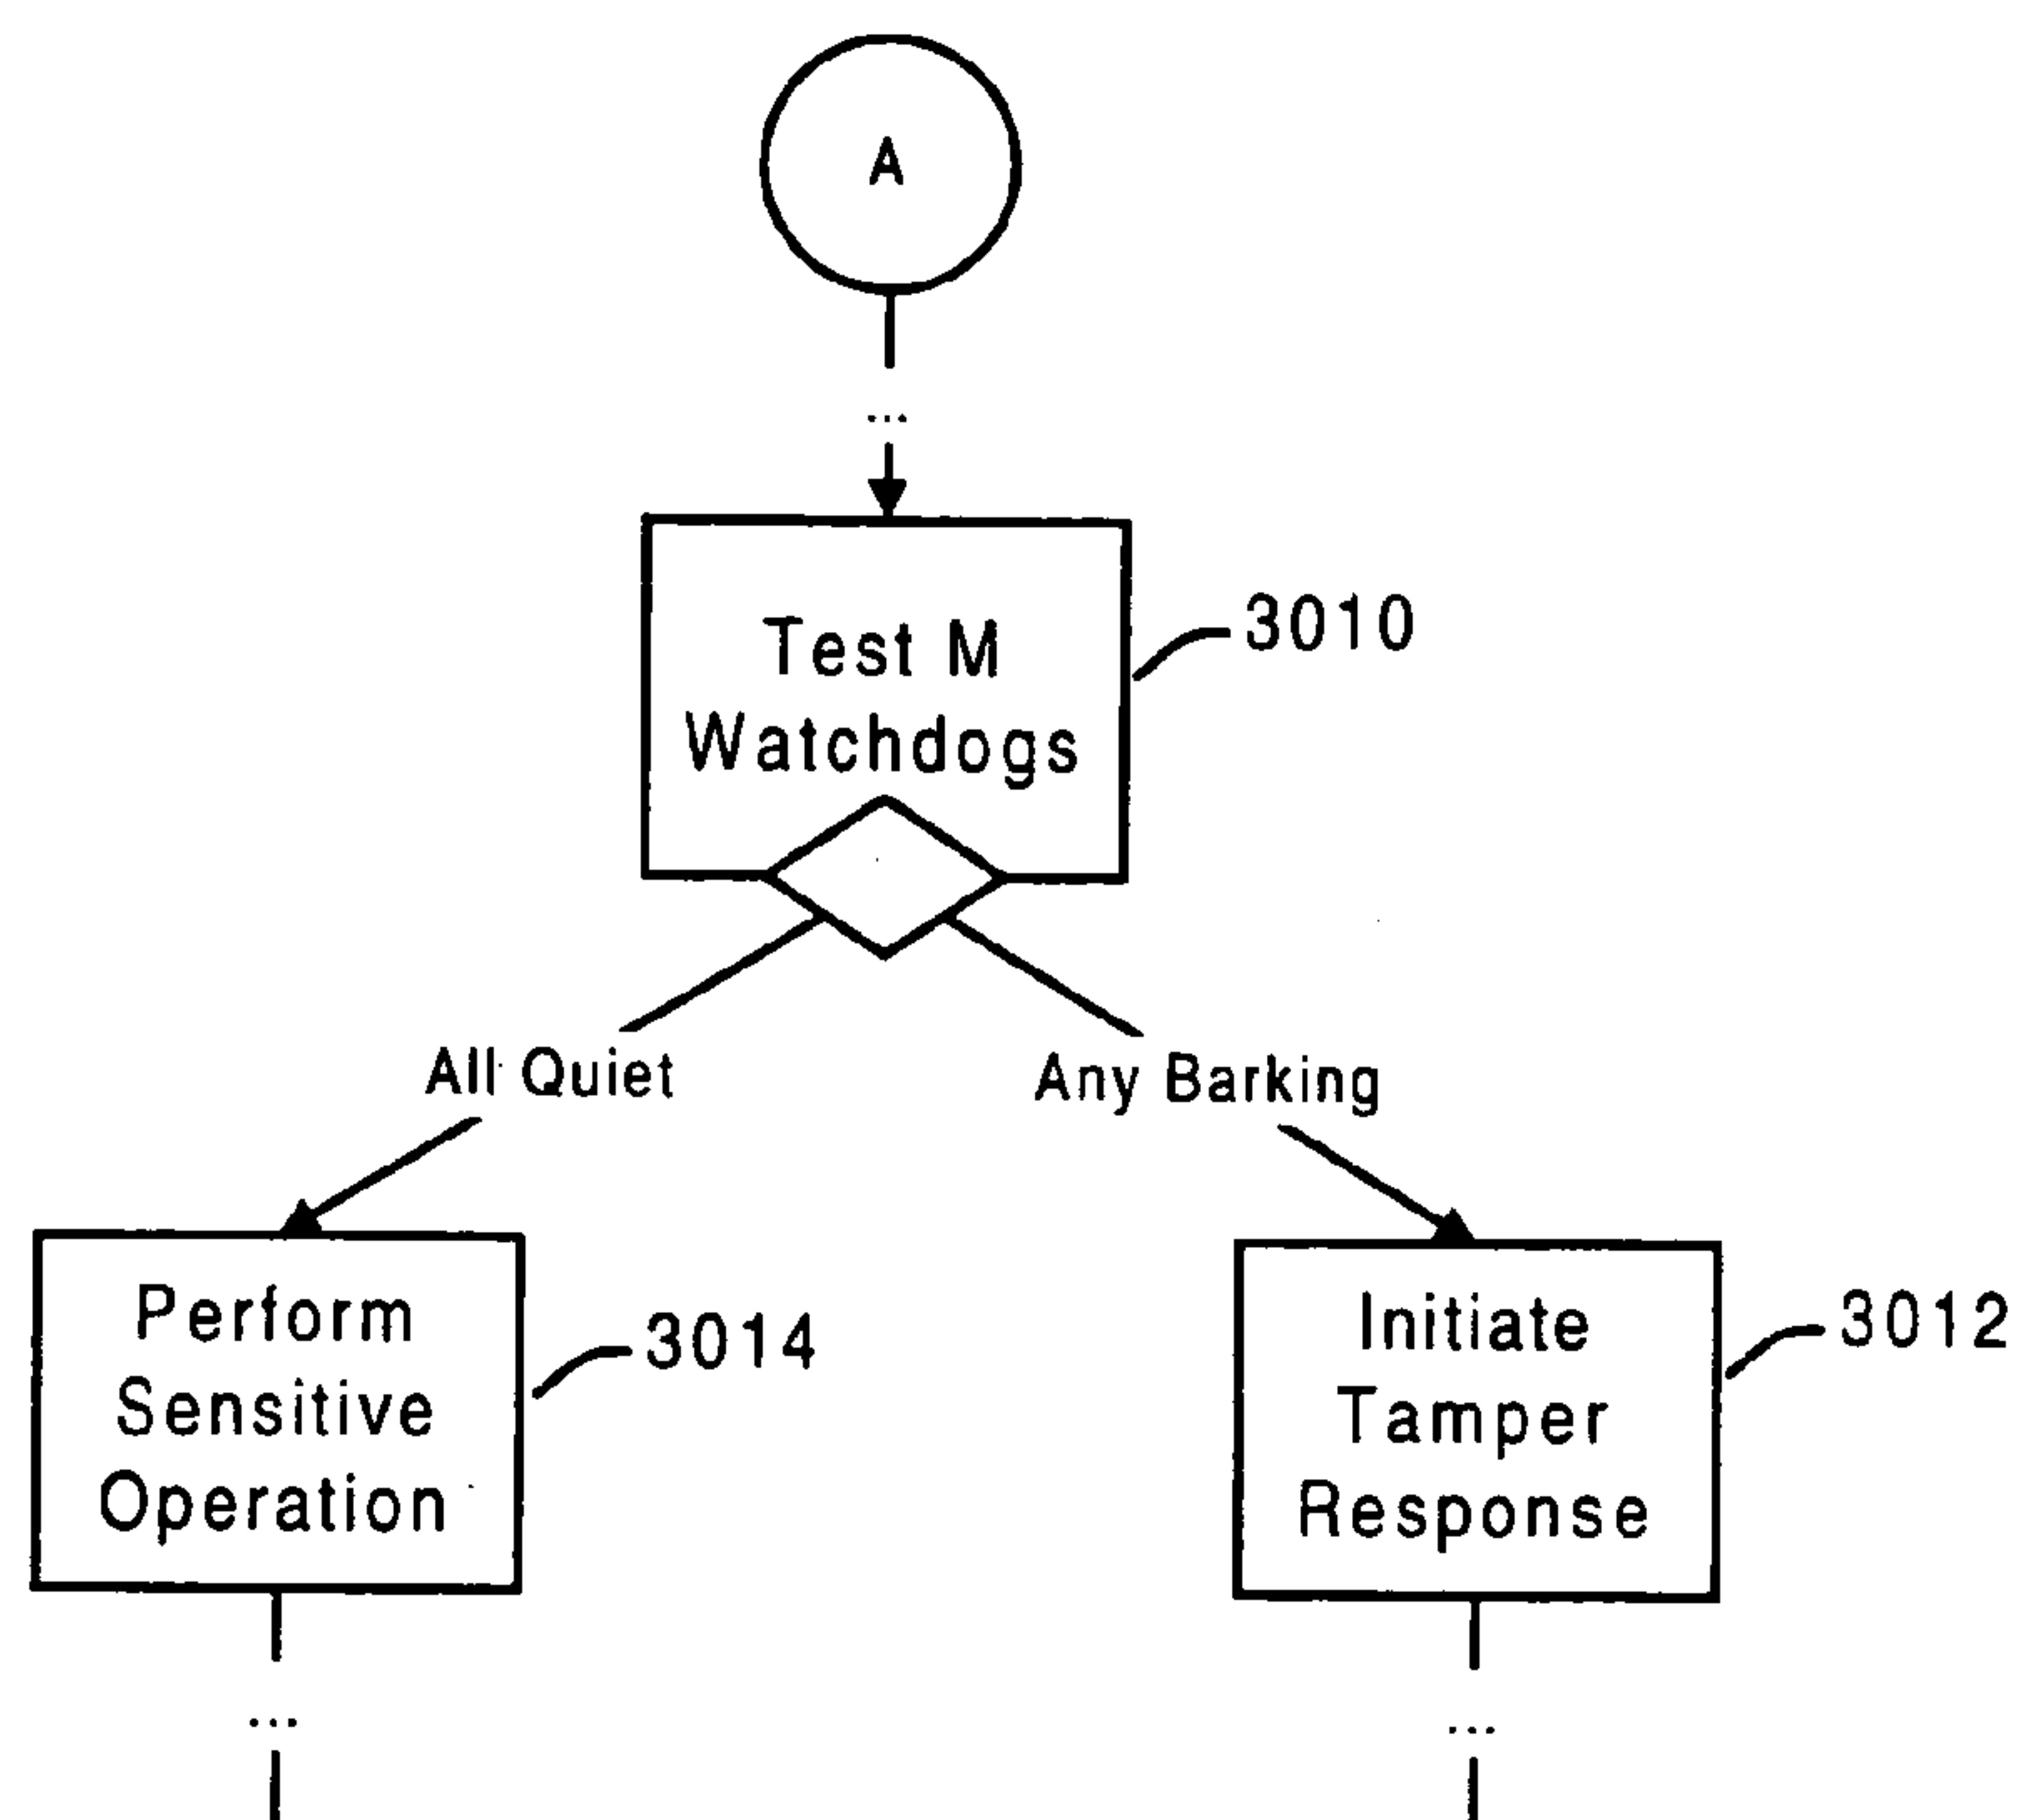


Fig. 30B

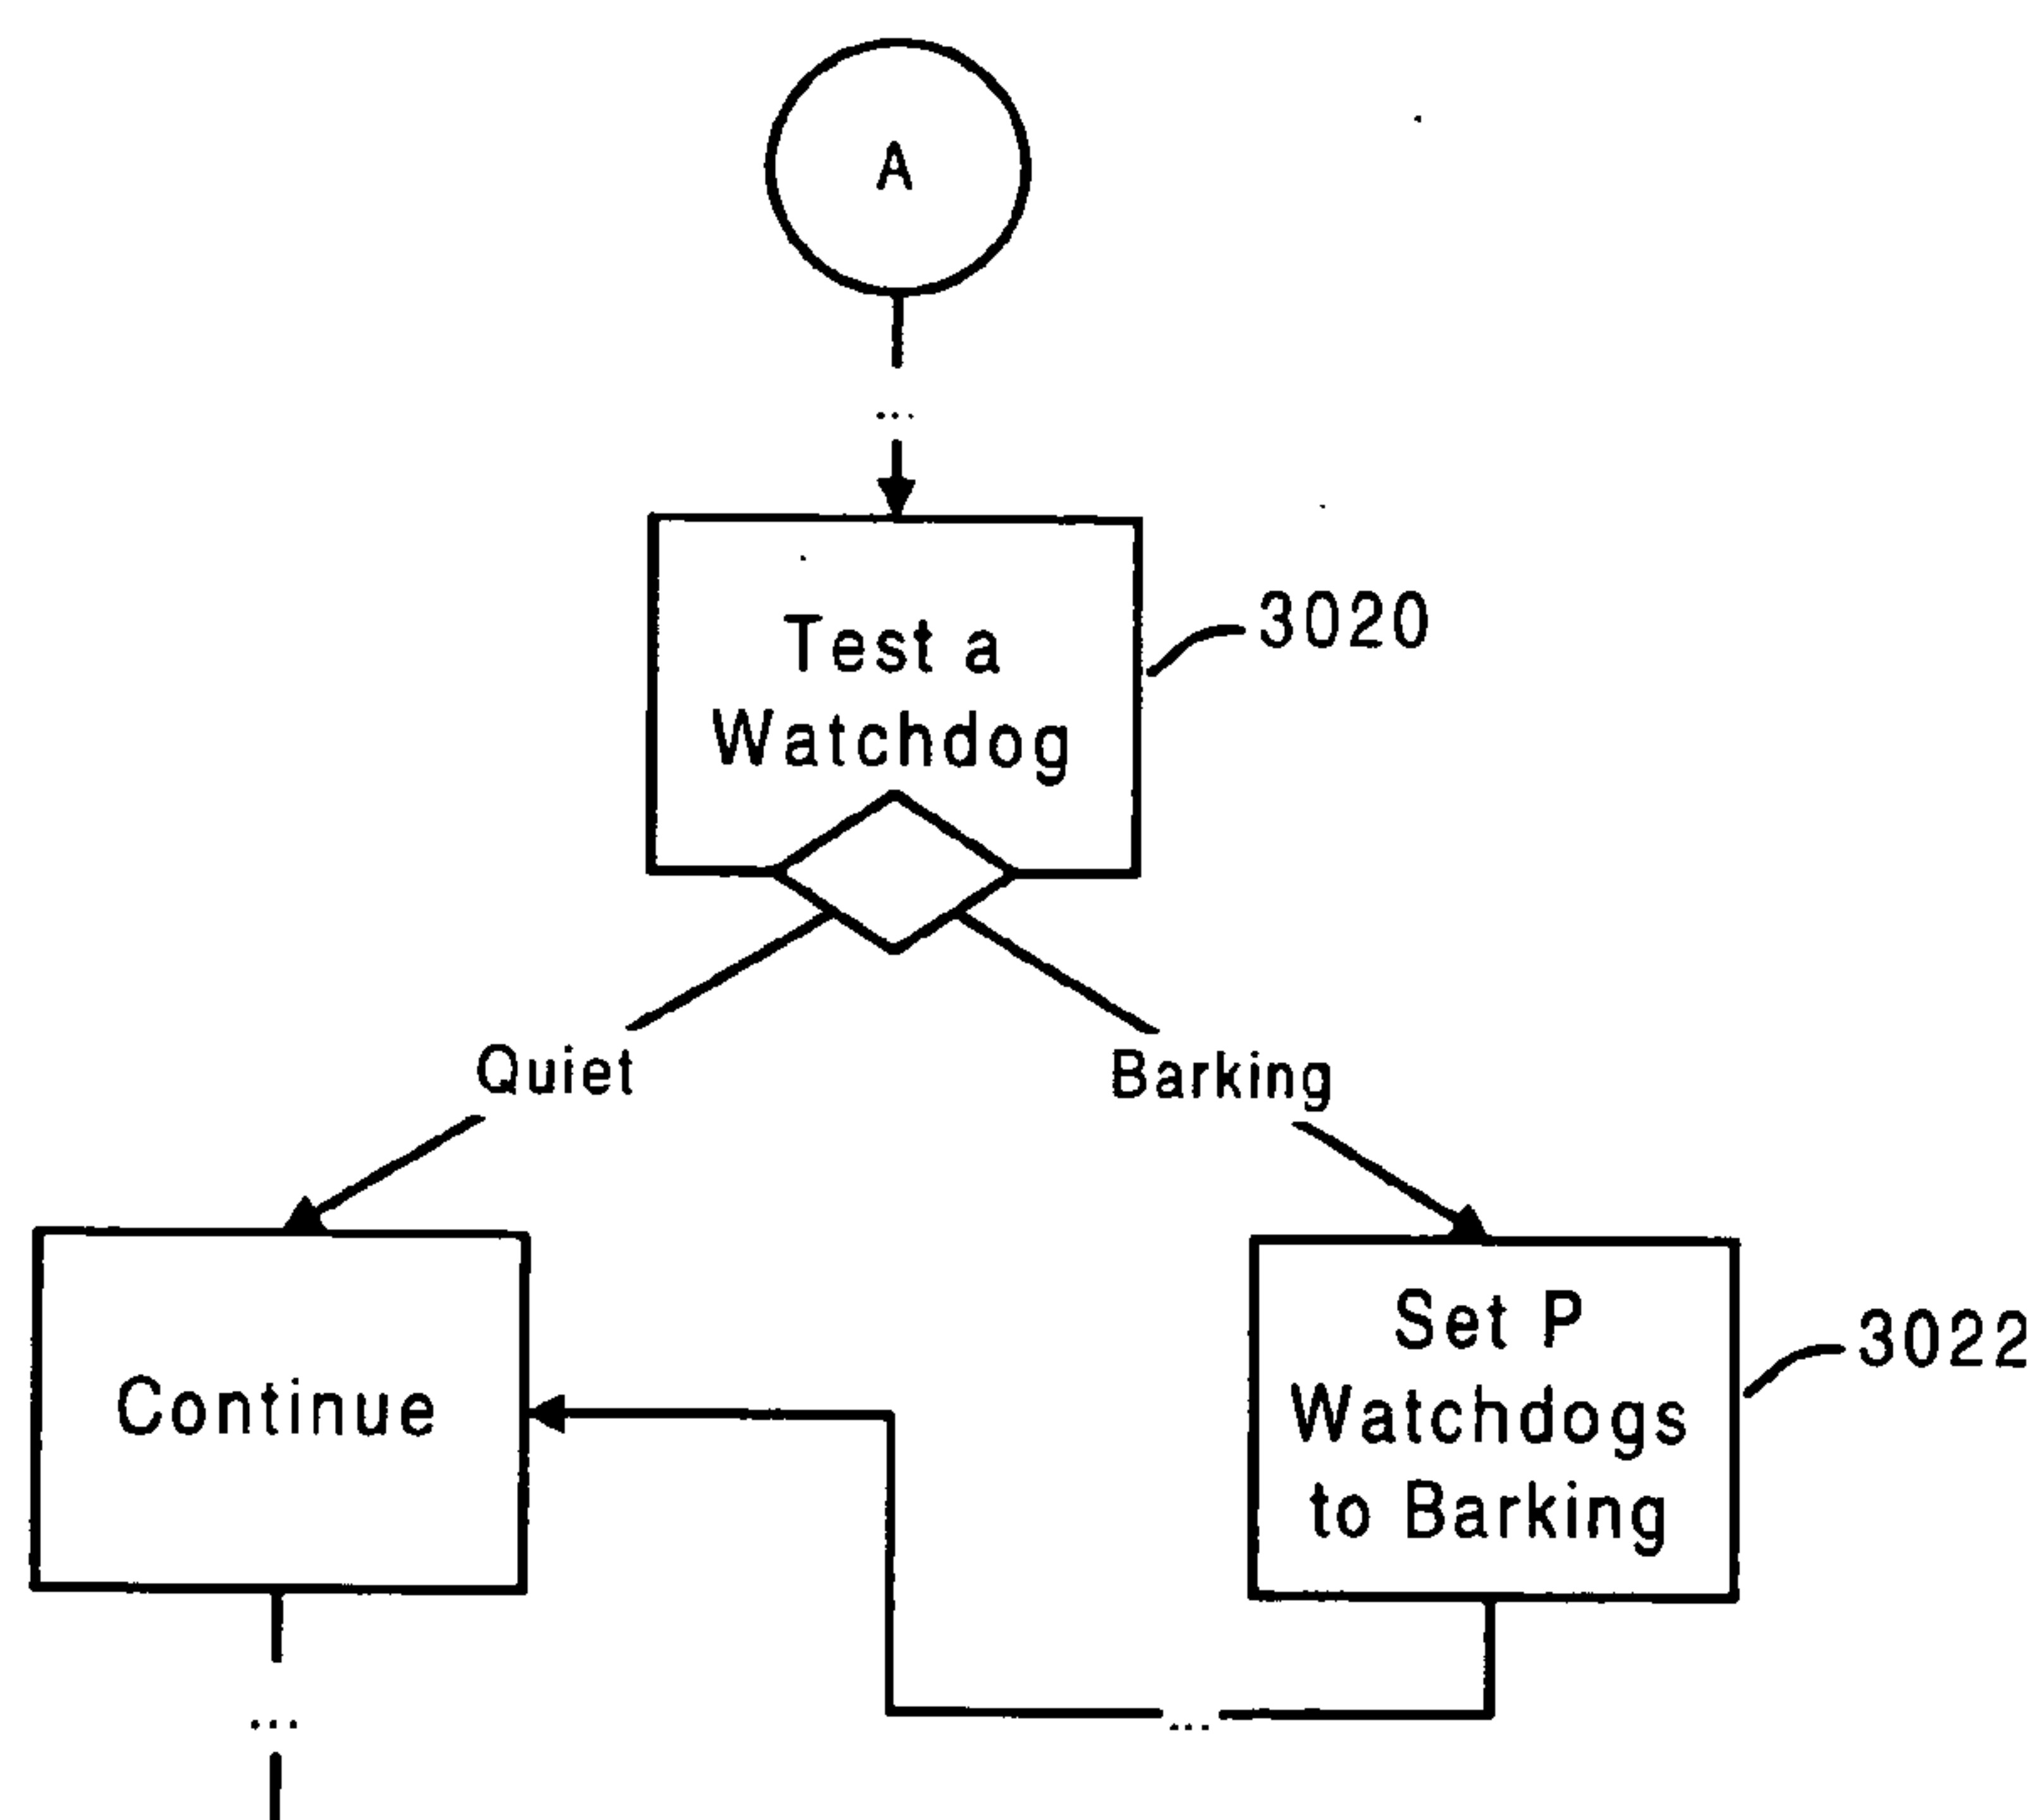


Fig. 30C




Fig. 31A




Fig. 31B

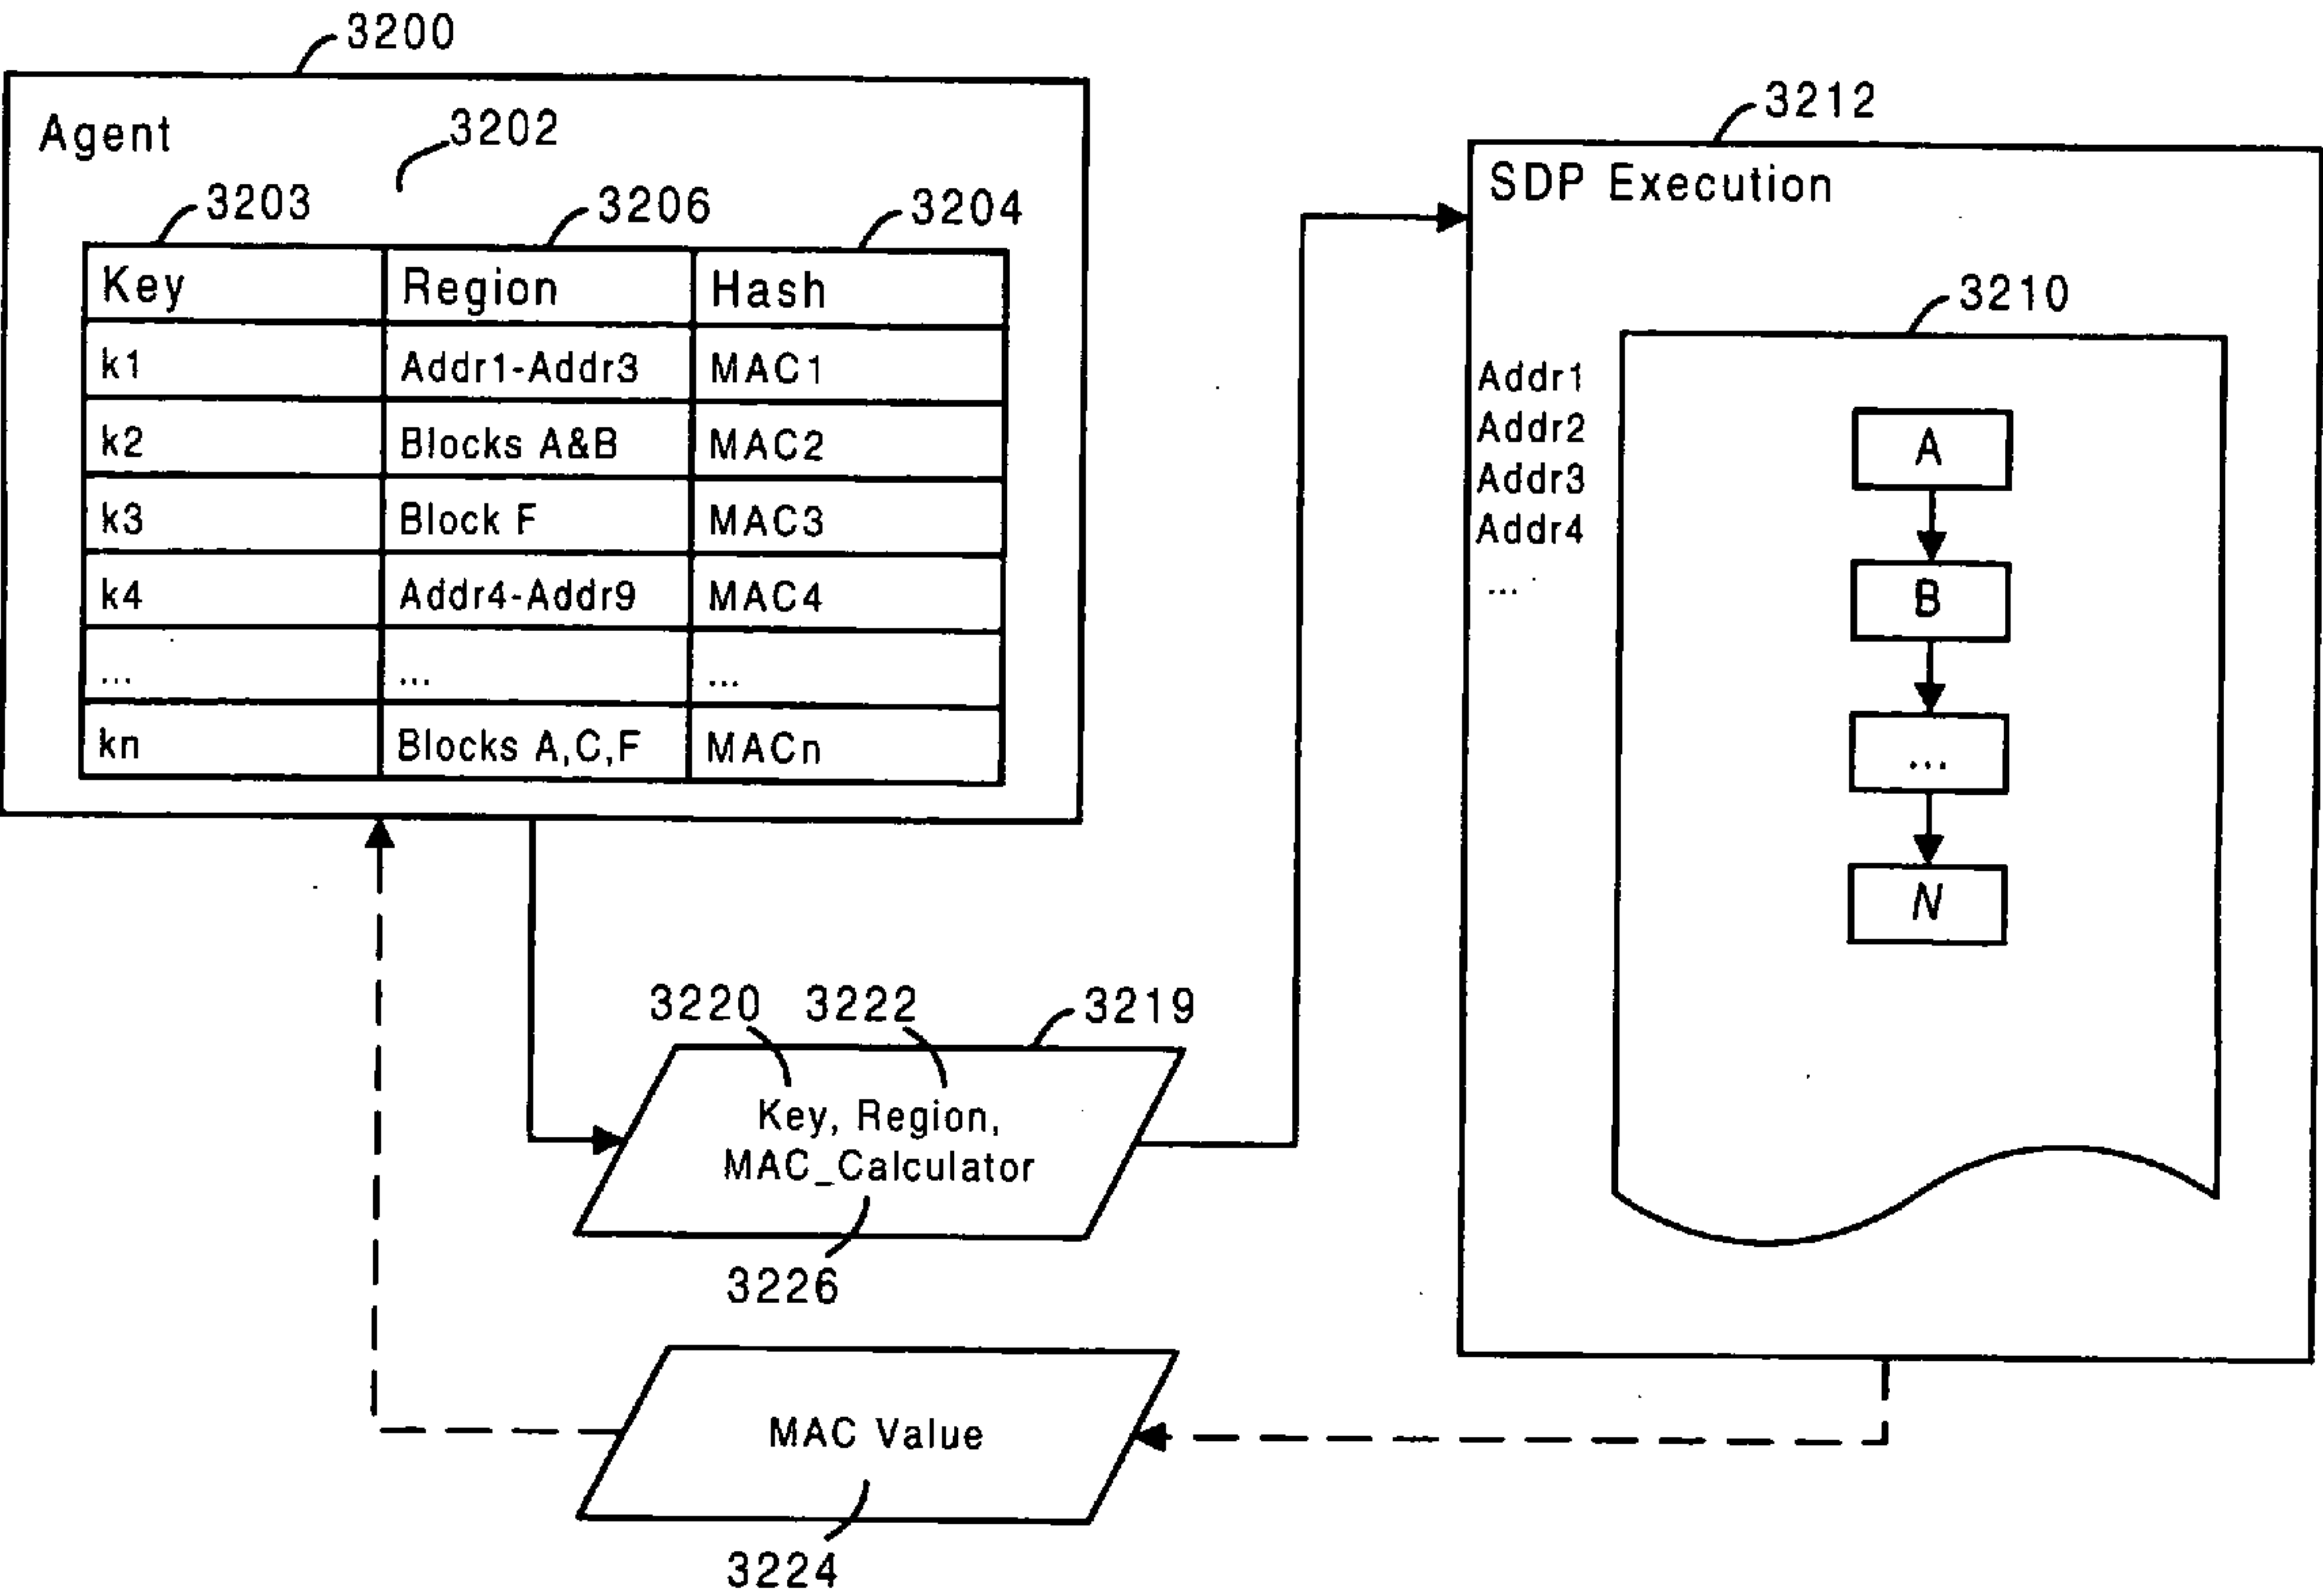


Fig. 32

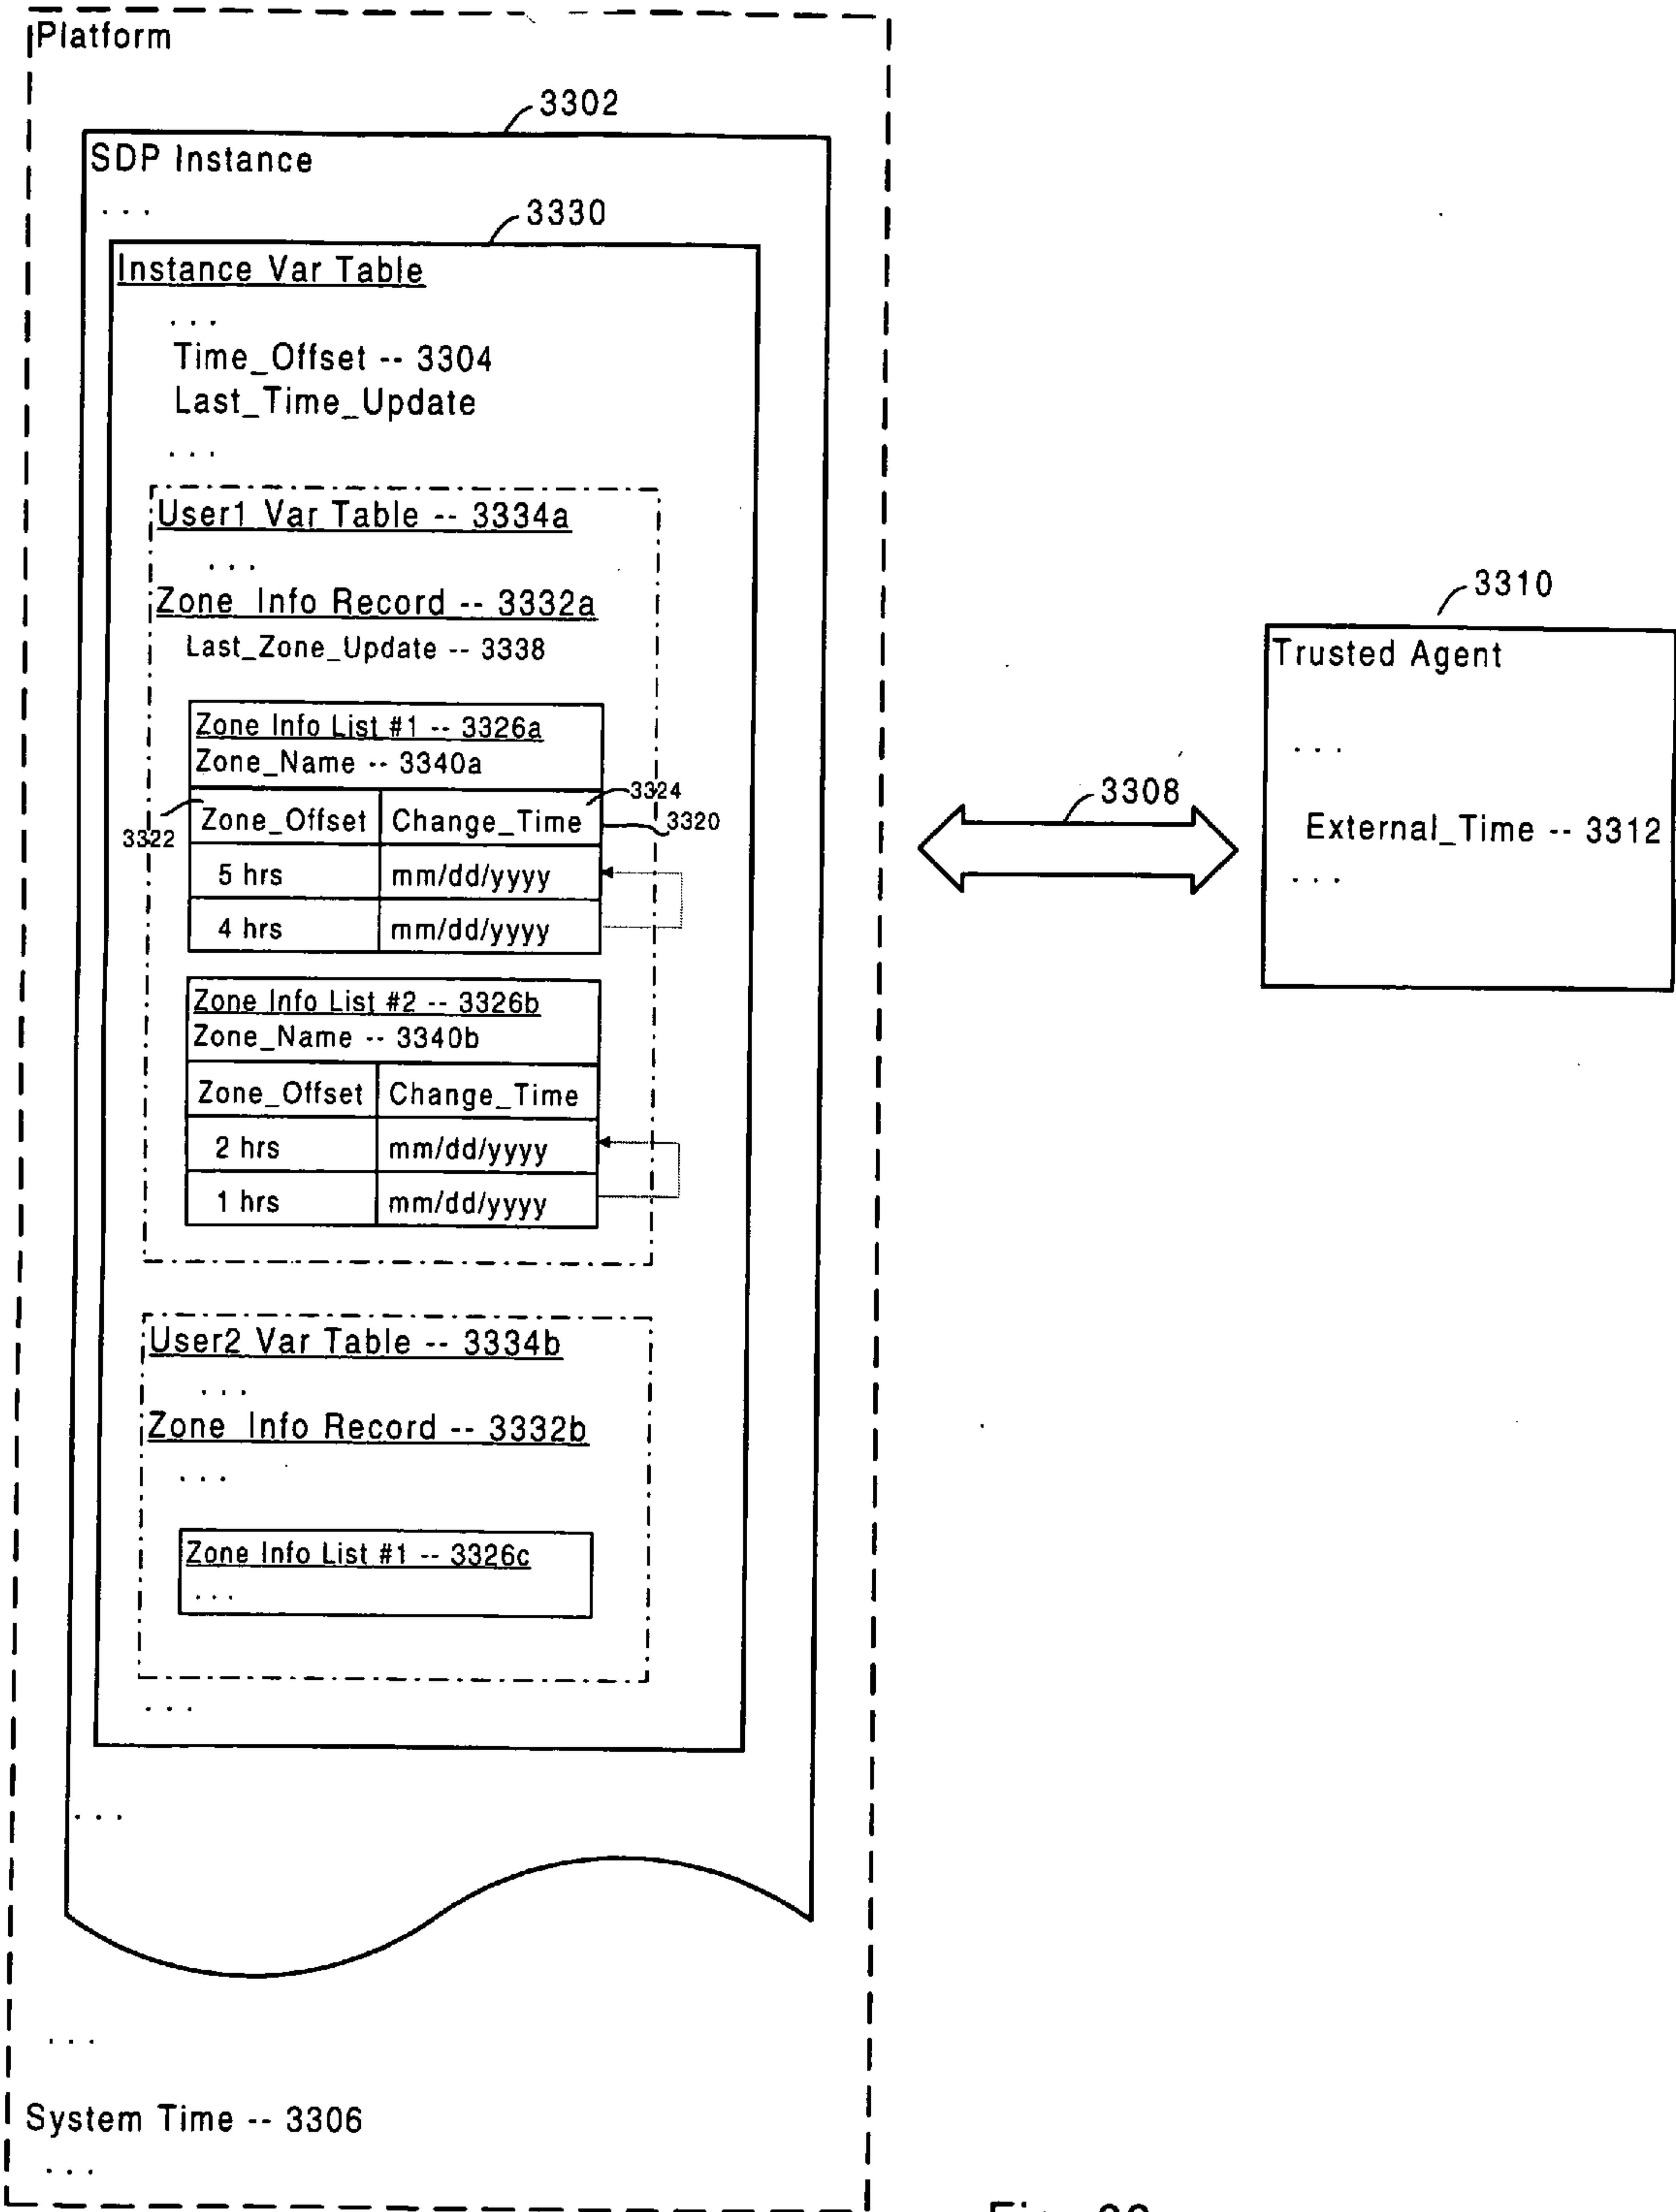


Fig. 33

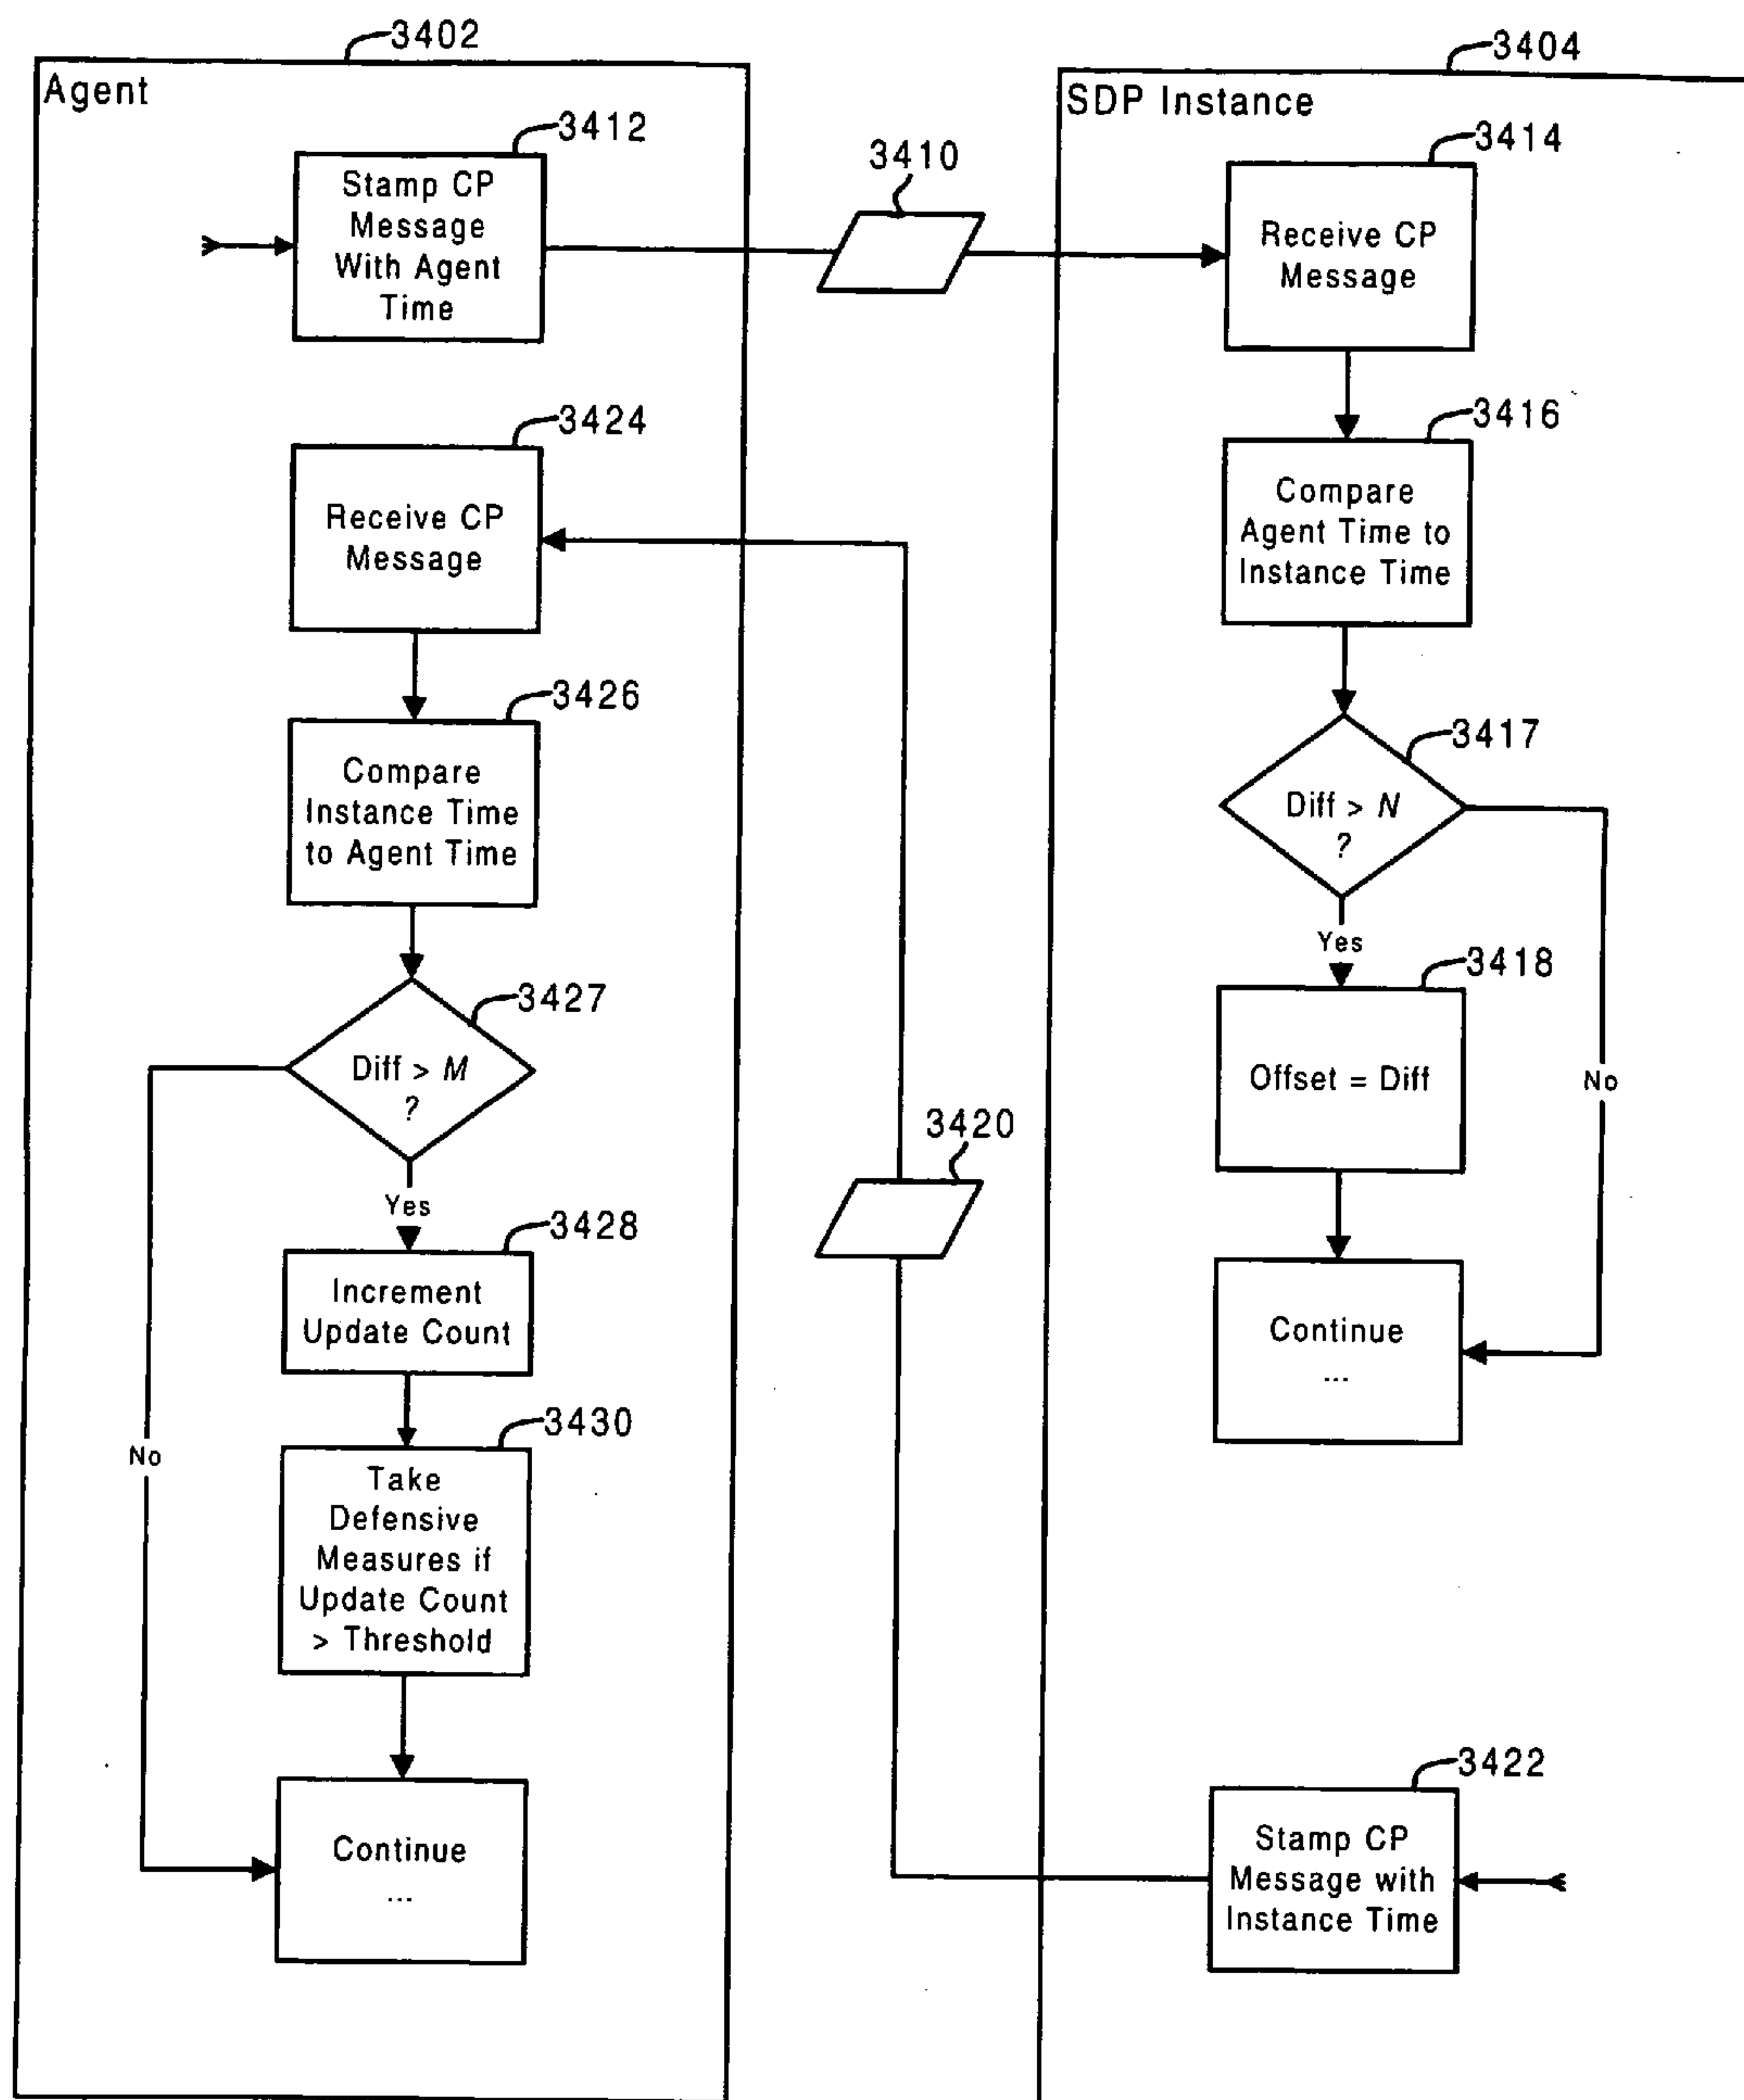


Fig. 34

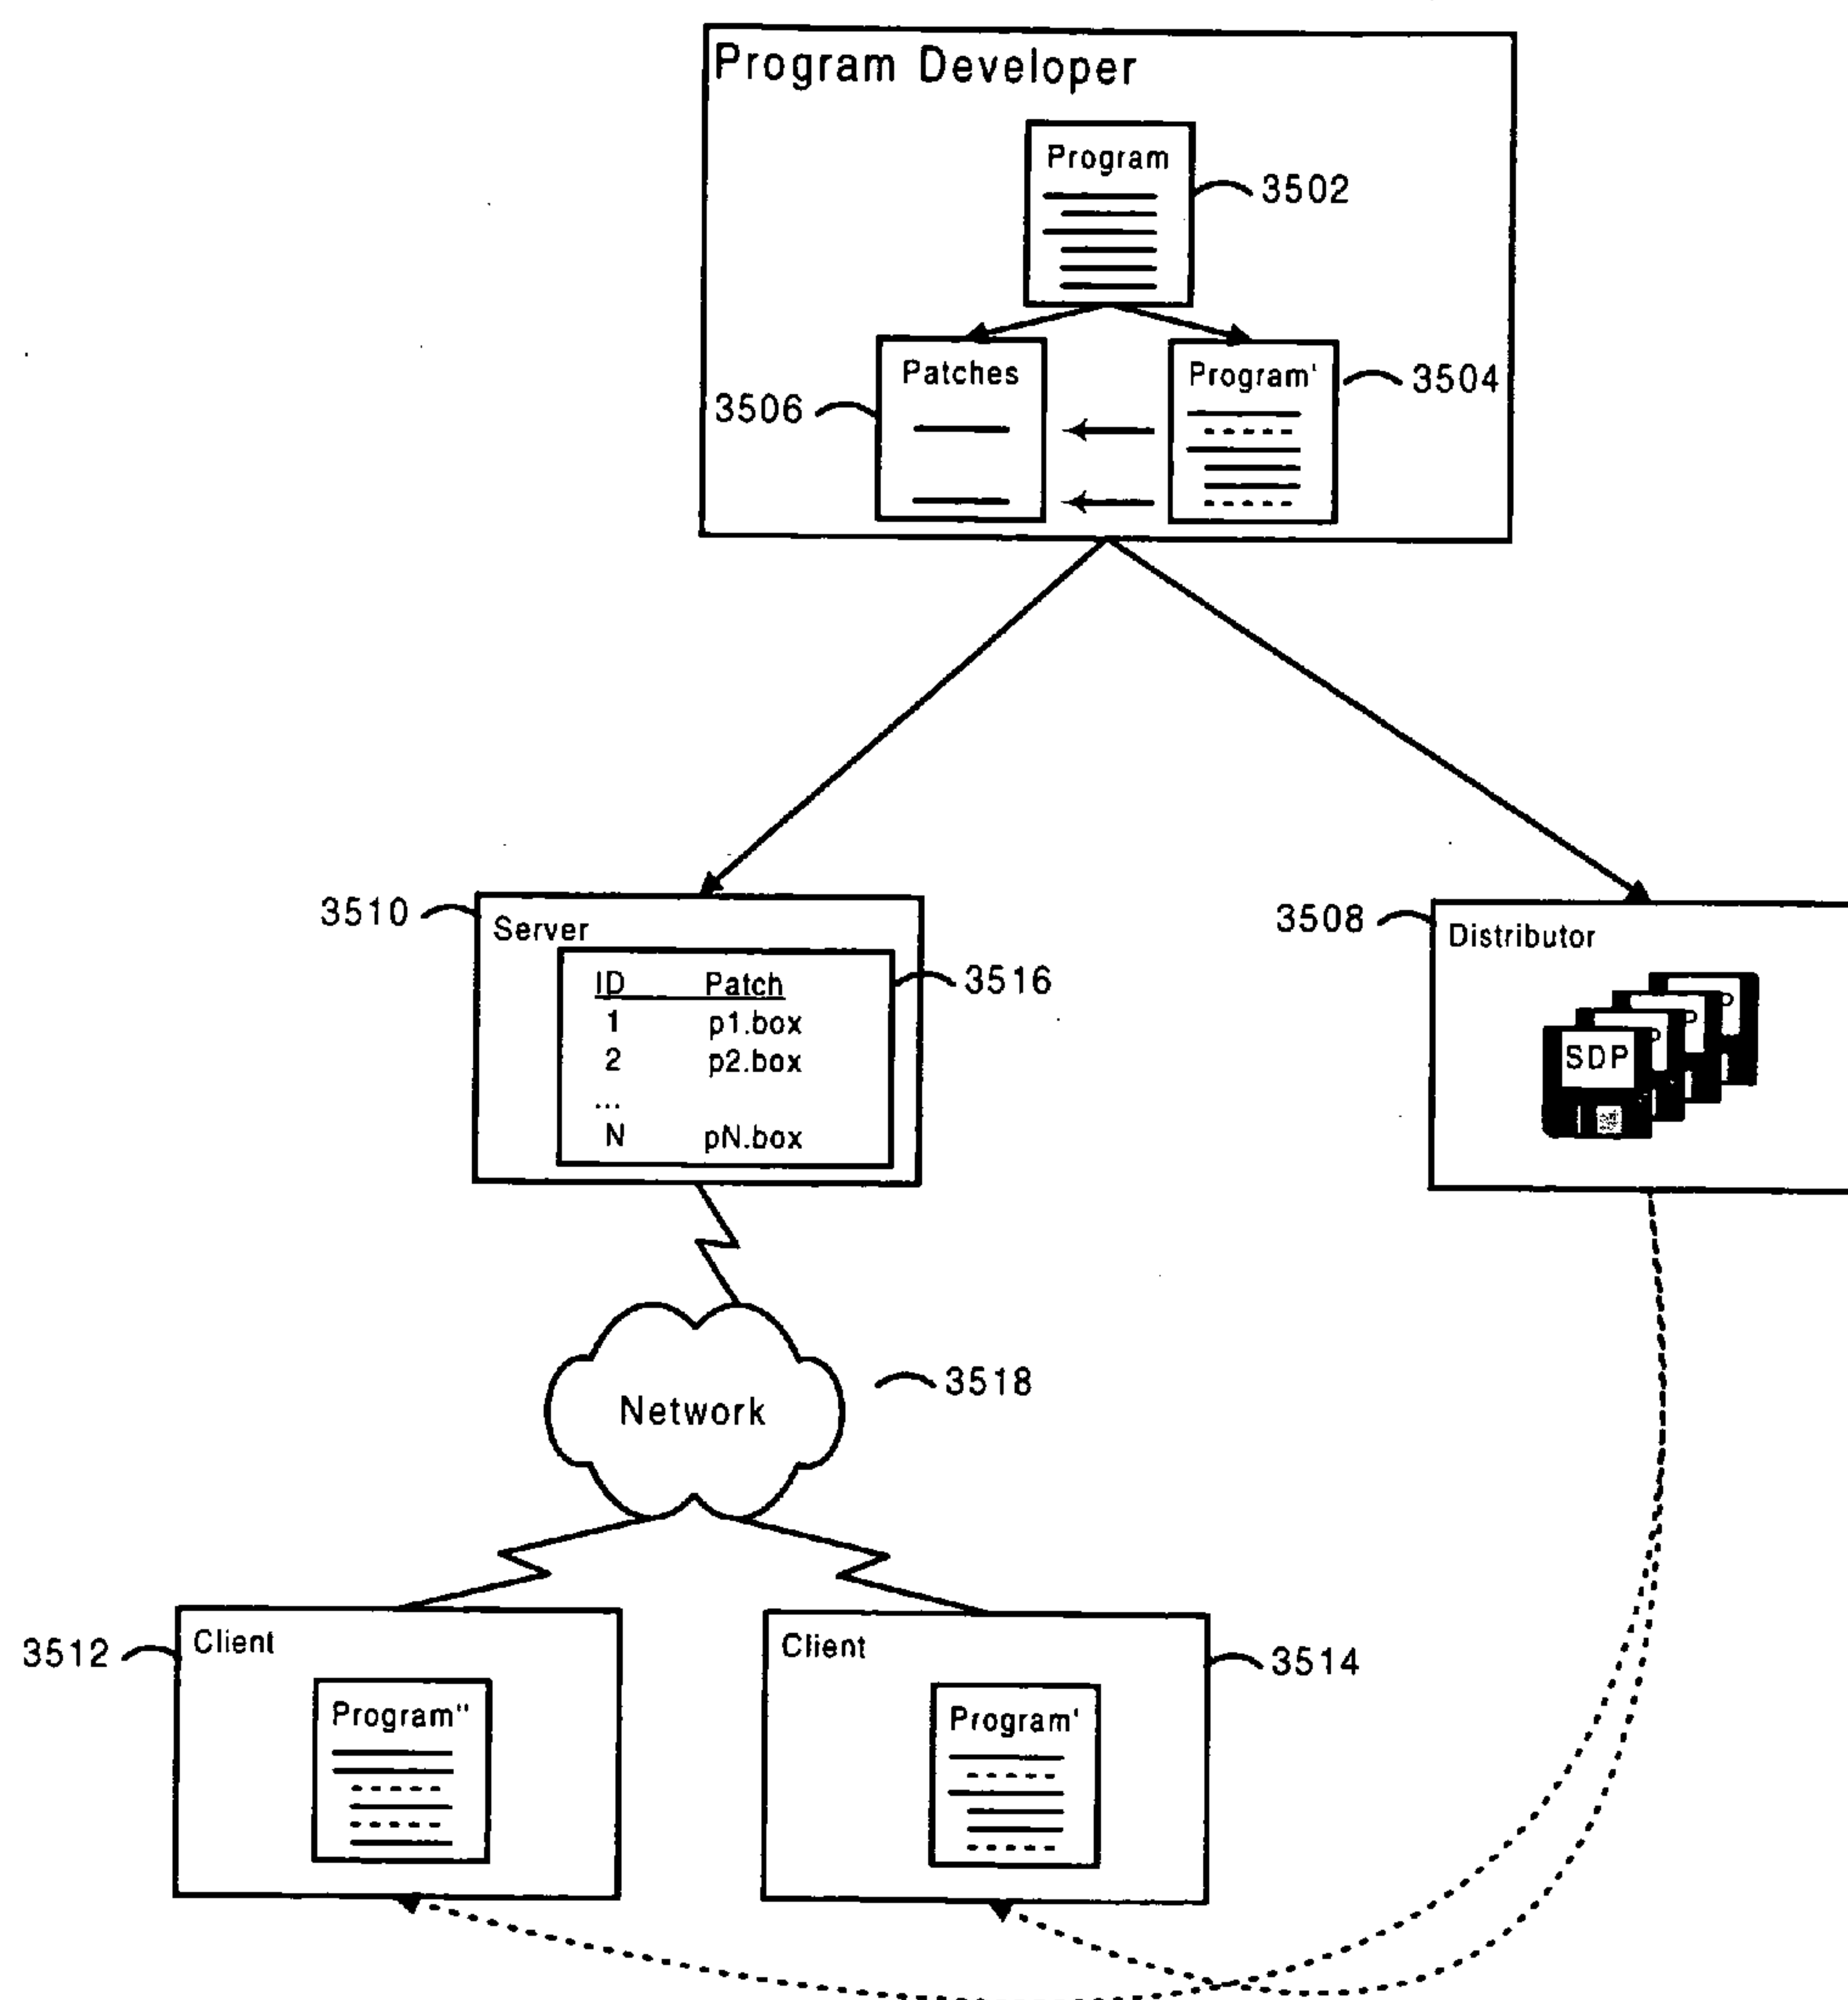


Fig. 35

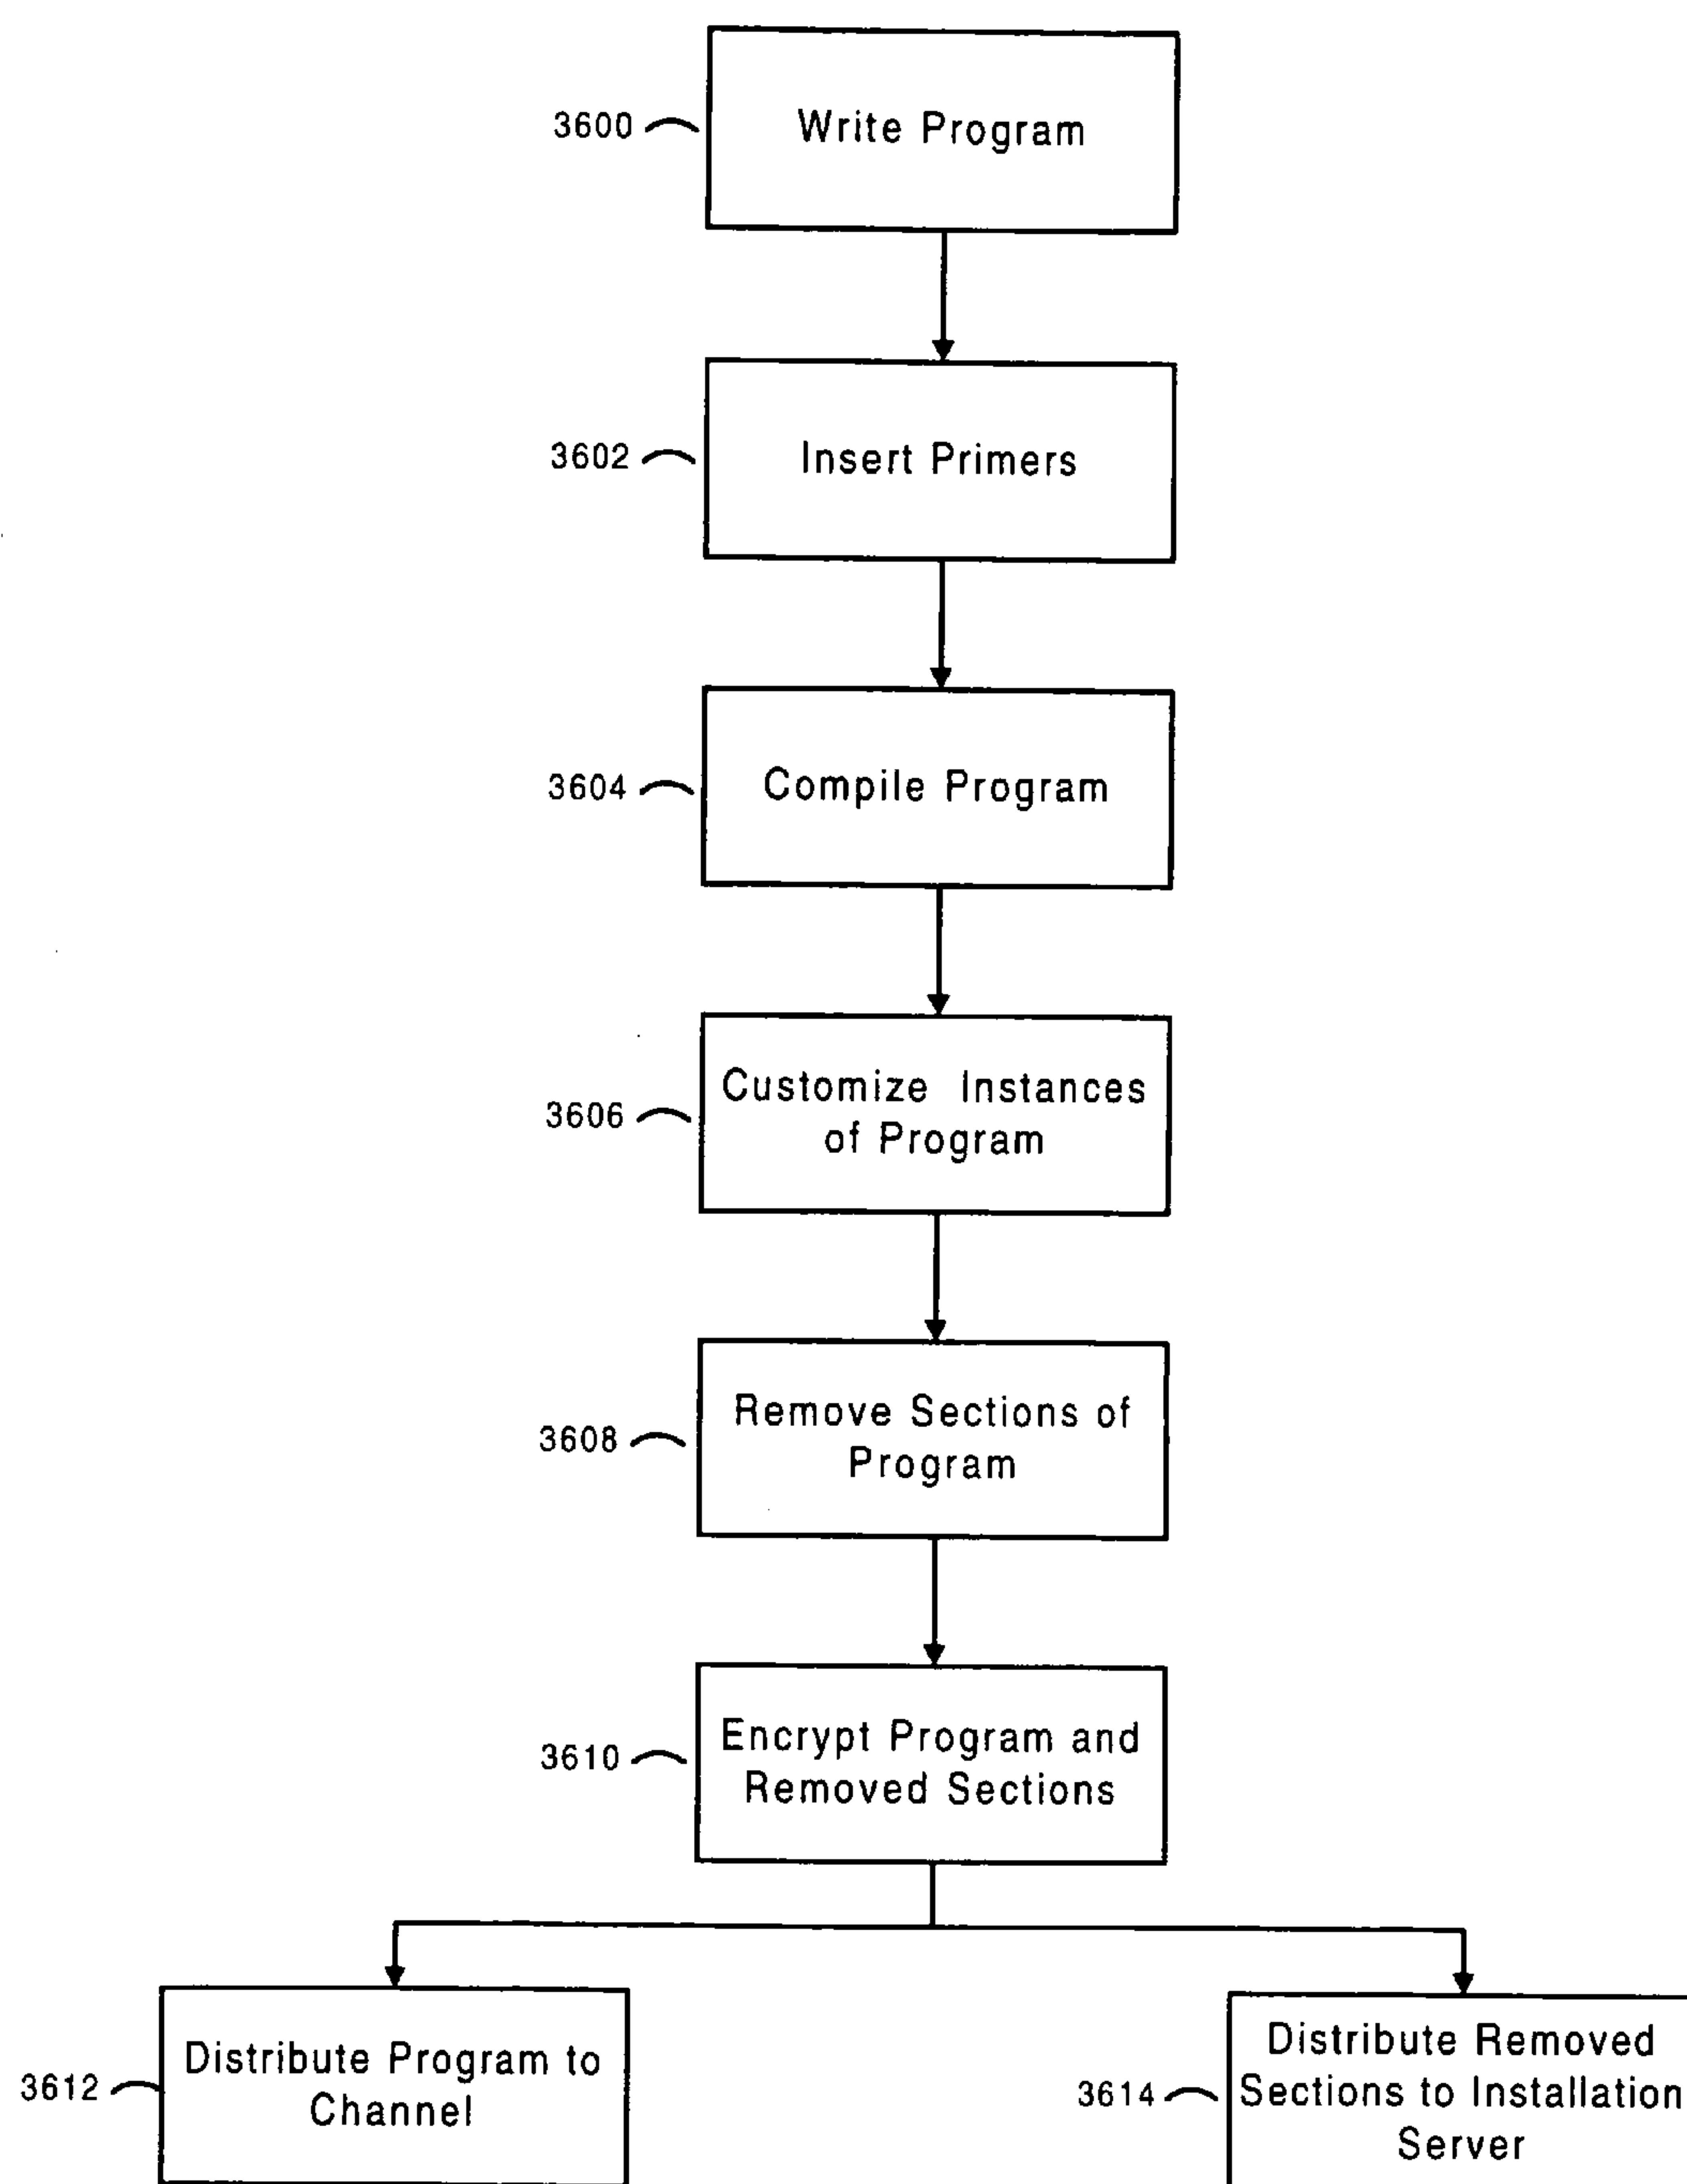


Fig. 36




Fig. 37A

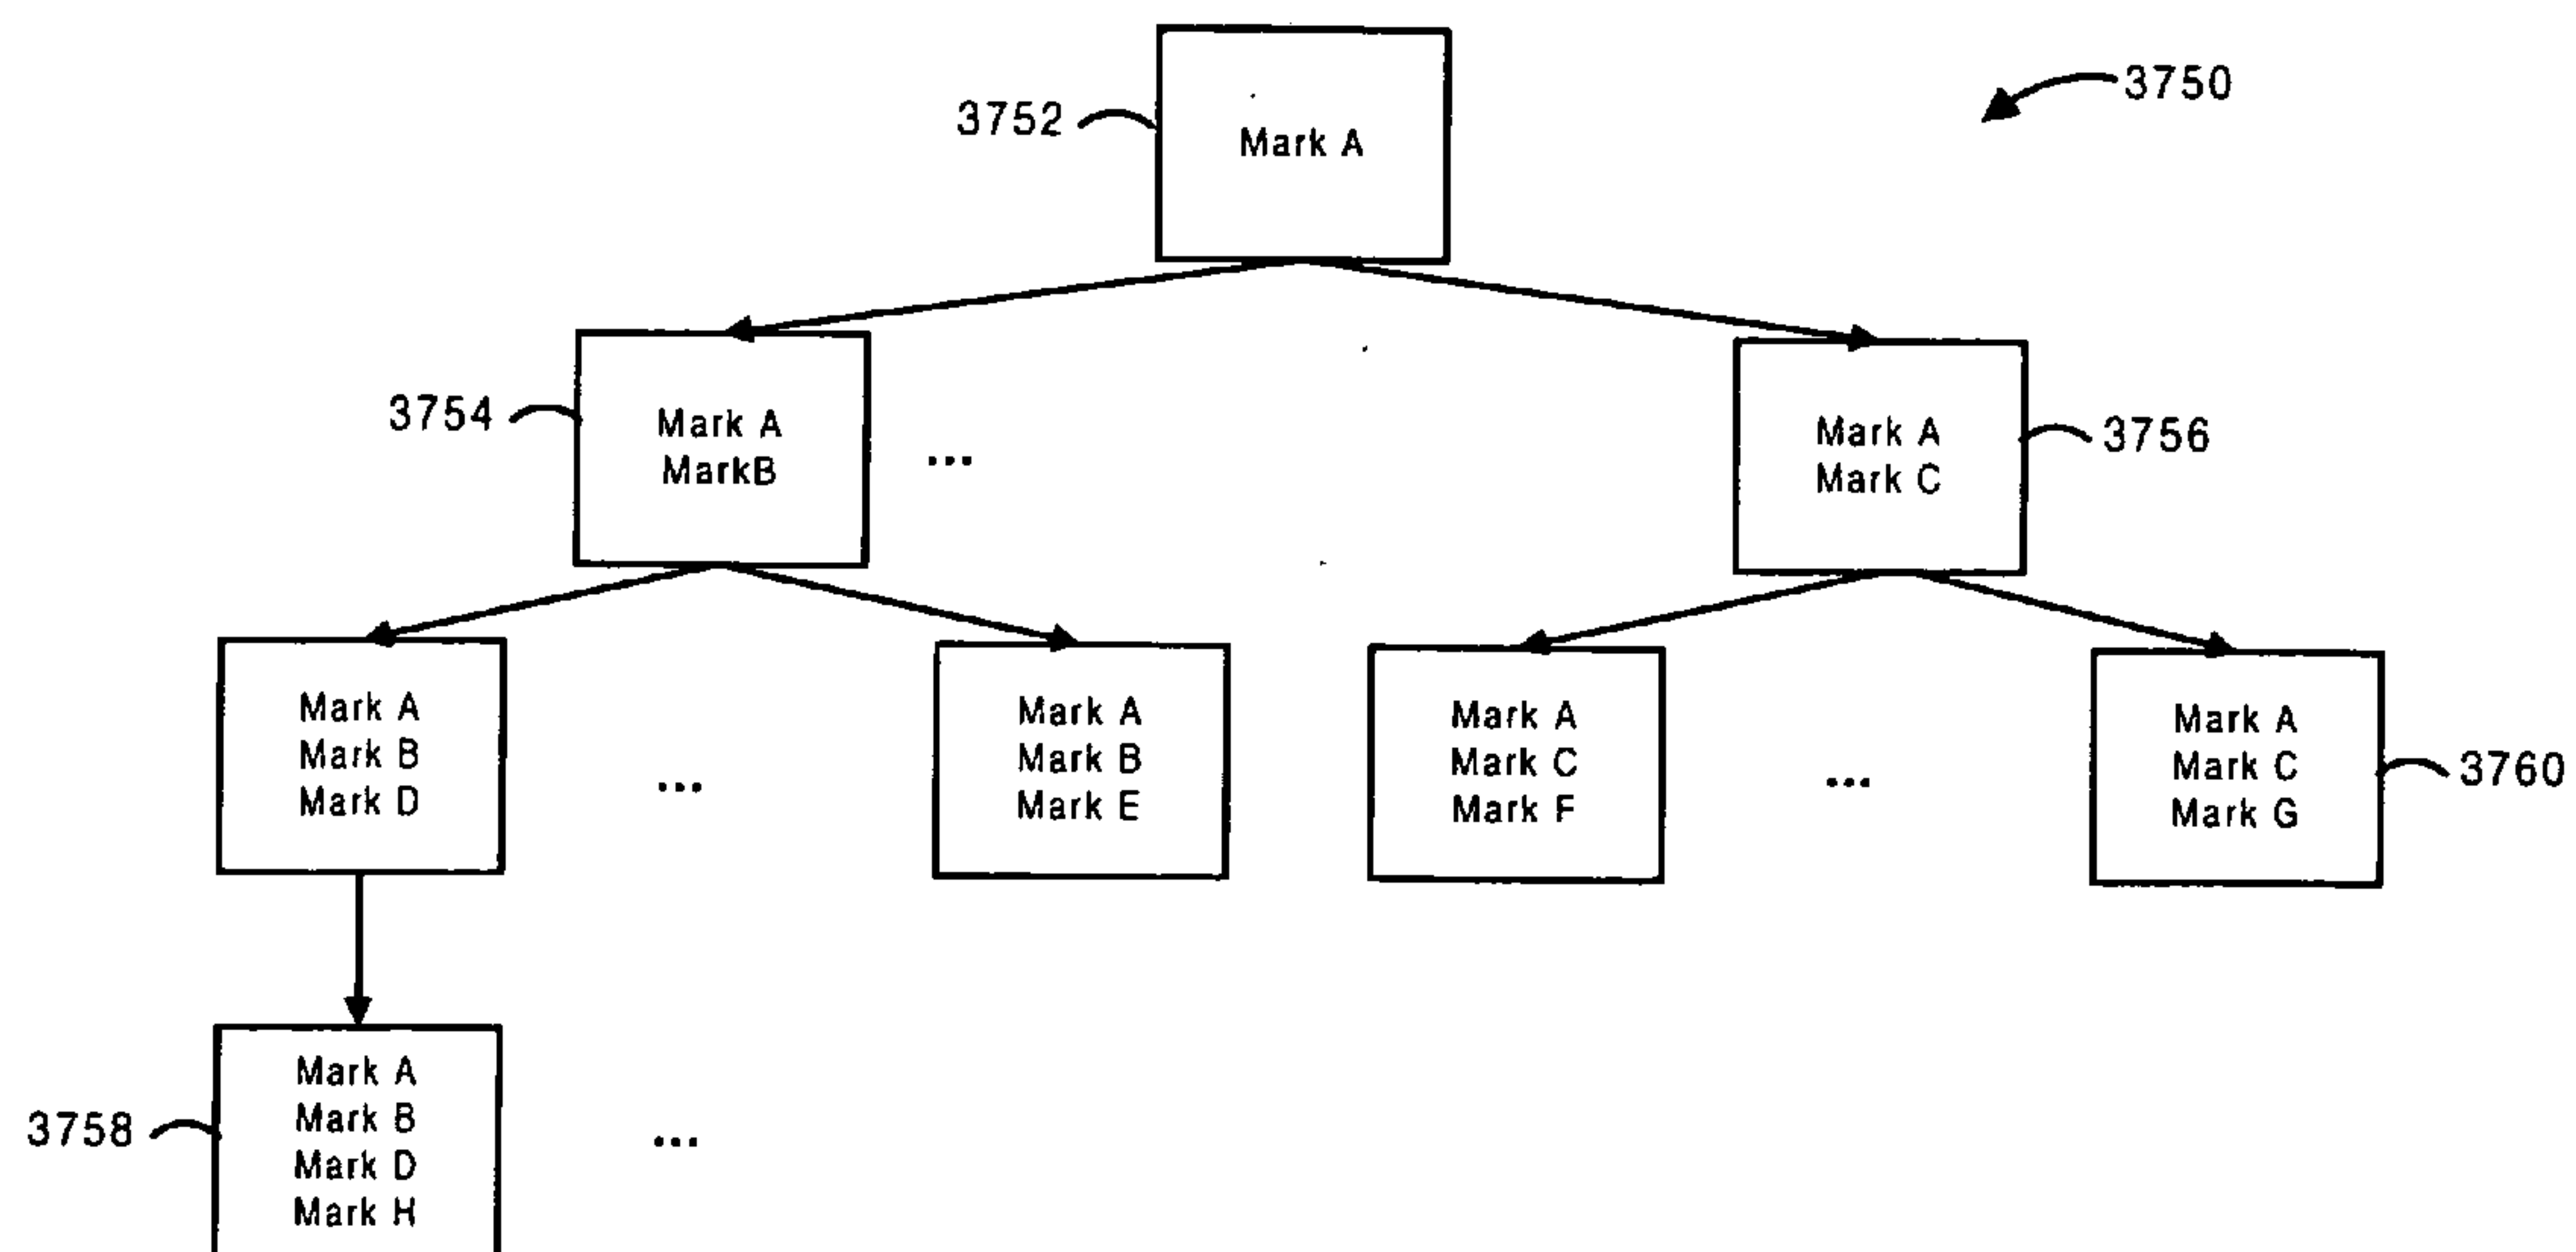


Fig. 37B

SOFTWARE SELF-DEFENSE SYSTEMS AND METHODS

RELATED APPLICATIONS

[0001] This application is a continuation of application Ser. No. 09/629,546, filed Jul. 31, 2000, which claims priority from U.S. Provisional Application Ser. No. 60/146,419, entitled "Software Self-Defense Systems and Methods," filed Jul. 29, 1999, and is related to commonly-assigned U.S. application Ser. No. 09/629,807, entitled "Systems and Methods for Watermarking Software and Other Media," filed Jul. 31, 2000 ("the Software Watermarking application"), all of which are hereby incorporated in their entirety by reference.

COPYRIGHT AUTHORIZATION

[0002] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

FIELD OF THE INVENTION

[0003] The present invention relates to the enhancement of software security. More particularly, the present invention relates to systems and methods for enabling software to detect and/or resist tampering, interpretation, decoding, reverse engineering, and/or other forms of attack.

BACKGROUND OF THE INVENTION

[0004] Advances in microelectronics and networking technology have enabled the development of increasingly complex computer systems. The software that controls these systems or operates thereon has also become exceedingly complex and costly to develop. As a result, companies will often wish to protect their software from reverse-engineering by competitors, from copying by pirates, or from tampering by persons attempting to undermine the integrity of the software's operation. While copyright law and patent law provide some level of protection, in many instances these legal protections are insufficient. What is needed are technical systems and methods for shielding the operational details of a piece of software from the view of others, and for resisting attempts to compromise the integrity of the software's operation. Yet due to the way software is often deployed, these are difficult tasks. For example, application software is often installed on a consumer's computer system, which is typically remote from the software developer's system. The consumer's system may thus provide an environment in which an attacker can analyze and modify the software with relative ease and with little risk of detection. Accordingly, systems and methods are also needed for protecting the secrecy and integrity of software when it is run in potentially hostile environments.

SUMMARY OF THE INVENTION

[0005] Systems and methods for protecting software programs and data from unauthorized analysis, attack, and/or modification are described herein. It should be appreciated that the present invention can be implemented in numerous

ways, including as a process, an apparatus, a system, a device, a method, or a computer readable medium. Several inventive embodiments of the present invention are described below.

[0006] In one embodiment, a system for obfuscating computer code is provided. The system includes a memory unit containing a binary code modification tool, a program module that includes a variety of routines for applying obfuscation transformations to computer code, and a computer program operable to associate or link the binary code modification tool with the program module to yield an obfuscation program. The system also includes a processor for executing the obfuscation program and other programs stored in the memory unit. When executed, the obfuscation program is operable to accept computer code as input, and to perform obfuscating transformations on the computer code, the obfuscating transformations tending to render the computer code more resistant to reverse engineering, decompilation, and/or other attacks.

[0007] In another embodiment, a method for obfuscating a computer program is described. The method includes the step of linking a binary code modification tool with a program module containing code for performing a variety of obfuscation transformations. The linked program is able to accept a target program as input, and to apply obfuscation transformations to the target program to yield an obfuscated version of the target program.

[0008] In yet another embodiment, a system is disclosed for protecting a computer program from unauthorized modification and analysis. The system includes means for obfuscating the computer program, means for incorporating one or more tamper-resistance mechanisms into the computer program, and means for watermarking the computer program. The computer program can be obfuscated by applying a variety of transformations that are designed to make the computer program more difficult to understand, modify, and/or reuse. The obfuscation transformations include techniques for obfuscating the computer program's local structure, control graph, and/or data structures, and include techniques that exploit concurrency and non-determinism. Tamper-resistance mechanisms can be used to determine whether the computer program is connected to a debugger or has code or data that have been modified. Robust records of whether tampering has been detected can be maintained. In addition, detected tampering can be reported to an external agent, and/or can trigger the intentional destruction or modification of sensitive information, or the termination of program execution. In one embodiment, a response to detected tampering is intentionally delayed from the time that tampering was detected.

[0009] In yet another embodiment, a collection of macros and a binary image postprocessor are used to add watermarks to the computer program, to implement tamper-resistance mechanisms, and/or to implement other self-defensive mechanisms. The macros are inserted into the computer program. The computer program may then be compiled or subjected to other transformations. The binary image postprocessor locates the distinctive patterns or "primers" that the macros create in the computer program. The binary image postprocessor then adds watermarks, code for performing tamper-resistance measures, and/or other code or data into the spaces reserved by the primers.

[0010] In another embodiment, the invention provides a technique for compensating for tampering with the clock on the platform on which a computer program is executed. A signed time offset is used to relate the platform's clock to a clock maintained by an external agent. The time offset is stored with the computer program, and when the computer program is executed the time offset is added to time measurements obtained from the platform's clock. The time offset is updated when the computer program communicates with the external agent.

[0011] These and other features and advantages of the present invention will be presented in more detail in the following detailed description and the accompanying figures which illustrate by way of example the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] The present invention will be readily understood by the following detailed description in conjunction with the accompanying drawings, wherein like reference numerals designate like structural elements, and in which:

[0013] FIG. 1 illustrates a system for practicing an embodiment of the present invention.

[0014] FIG. 2 illustrates the functionality and operation of a binary program modification tool in accordance with one embodiment of the present invention.

[0015] FIG. 3 illustrates a block rearrangement technique in accordance with the principles of the present invention.

[0016] FIG. 4 illustrates a block splitting technique in accordance with the principles of the present invention.

[0017] FIG. 5 is an illustration of a block-duplication technique in accordance with an embodiment of the present invention.

[0018] FIG. 6 illustrates the addition of null-effect sequences and instructions to a program.

[0019] FIGS. 7A and 7B illustrate a stack-execution technique in accordance with the principles of the present invention.

[0020] FIGS. 8A, 8B, 8C, 8D, and 8E illustrate a module for creating spaces in executable code and for making them available for data and/or instructions.

[0021] FIG. 9 illustrates a technique for interfering with a debugger.

[0022] FIG. 10 illustrates a technique for detecting the presence of a debugger in accordance with the principles of the present invention.

[0023] FIG. 11 illustrates a technique for interfering with the operation of debugger by storing instructions and/or data into the instruction stream.

[0024] FIG. 12 illustrates a code encryption technique in accordance with the principles of the present invention.

[0025] FIG. 13 illustrates a code self-modification technique in accordance with the principles of the present invention.

[0026] FIG. 14 illustrates a technique for obscuring external calls.

[0027] FIG. 15 illustrates a control graph obfuscation technique in accordance with an embodiment of the present invention.

[0028] FIGS. 16A, 16B, 16C, and 16D illustrate code replication techniques in accordance with the principles of the present invention.

[0029] FIGS. 17A and 17B illustrate a control graph obfuscation technique in accordance with the principles of the present invention.

[0030] FIG. 18 illustrates a code fusion technique in accordance with an embodiment of the present invention.

[0031] FIG. 19 illustrates an additional embodiment of a code fusion technique in accordance with an embodiment of the present invention.

[0032] FIG. 20 illustrates the use of overlays.

[0033] FIG. 21 illustrates a technique for relocating code in accordance with the principles of the present invention.

[0034] FIGS. 22A, 22B, 22C, and 22D illustrate a technique for protecting jumps or other conditional code sequences from attack.

[0035] FIG. 23 illustrates the use of non-determinism in accordance with an embodiment of the present invention.

[0036] FIG. 24 illustrates the use of concurrency to obfuscate the operation of a program at start-up.

[0037] FIG. 25 illustrates a data scrambling technique in accordance with an embodiment of the present invention.

[0038] FIG. 26 illustrates a priming and stamping technique in accordance with an embodiment of the present invention.

[0039] FIG. 27 illustrates the use of a self-validation technique in accordance with an embodiment of the present invention.

[0040] FIGS. 28A and 28B illustrate a mechanism for applying a self-validation technique in accordance with an embodiment of the present invention.

[0041] FIGS. 29A, 29B, and 29C illustrate a technique for preventing an attacker from repeatedly restarting a self-defensive program.

[0042] FIGS. 30A, 30B, 30C illustrate the use of watchdog bits in one embodiment of the present invention.

[0043] FIGS. 31A and 31B illustrate tamper reaction mechanisms in accordance with embodiments of the present invention.

[0044] FIG. 32 is an illustration of a code validation technique in accordance with an embodiment of the present invention.

[0045] FIG. 33 illustrates a technique for maintaining secure time and time zone information in accordance with an embodiment of the present invention.

[0046] FIG. 34 illustrates a technique for resisting tampering with the system time as seen by a self-defensive program in one embodiment of the present invention.

[0047] FIG. 35 illustrates a system and method for applying one or more watermarks to a program in accordance with an embodiment of the present invention.

[0048] FIG. 36 is a flow chart of a method for applying one or more watermarks to a program in accordance with an embodiment of the present invention.

[0049] FIGS. 37A and 37B illustrate a watermark application technique in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[0050] A detailed description of the present invention is provided below. While the invention is described in conjunction with several embodiments, it should be understood that the invention is not limited to any one embodiment. On the contrary, the scope of the invention is limited only by the appended claims, and the invention encompasses numerous alternatives, modifications and equivalents. Although numerous specific details are set forth in the following description in order to provide a thorough understanding of the present invention, it should be understood that the present invention may be practiced according to the claims without some or all of these details. Moreover, for purposes of clarity, certain technical material that is known in the fields related to the invention has not been described in detail in order to avoid unnecessarily obscuring the invention.

[0051] The following U.S. patents and applications, each of which is assigned to the assignee of the current application, are hereby incorporated in their entirety by reference: Ginter et al., “Systems and Methods for Secure Transaction Management and Electronic Rights Protection,” U.S. Pat. No. 5,915,019, issued Jun. 22, 1999 (“the Ginter ’019 patent”); Ginter et al., “Trusted Infrastructure Support Systems, Methods and Techniques for Secure Electronic Commerce, Electronic Transactions, Commerce Process Control Automation, Distributed Computing, and Rights Management,” U.S. patent application Ser. No. 08/699,712, filed Aug. 12, 1996 (“the Ginter ’712 application”); Van Wie et al., “Steganographic Techniques for Securely Delivering Electronic Digital Rights Management Control Information Over Insecure Communications Channels,” U.S. Pat. No. 5,943,422, issued Aug. 24, 1999 (“the Van Wie ’422 patent”); Ginter et al., “Systems and Methods for Secure Transaction Management and Electronic Rights Protection,” U.S. Pat. No. 5,892,900, issued Apr. 6, 1999 (“the Ginter ’900 patent”); Shear et al., “Systems and Methods Using Cryptography to Protect Secure Computing Environments,” U.S. patent application Ser. No. 08/689,754, filed Aug. 12, 1996 (“the Shear ’754 application”); Maheshwari et al., “Trusted Storage Systems and Methods,” U.S. patent application Ser. No. 09/617,148, filed Jul. 17, 2000 (“the Maheshwari application”); Sibert, “Systems and Methods for Using Cryptography to Protect Secure and Insecure Computing Environments,” U.S. patent application Ser. No. 09/628,692, filed Jul. 28, 2000 (“the Sibert application”); Shear et al., “Cryptographic Methods, Apparatus & Systems for Storage Media Electronic Rights Management in Closed & Connected Appliances,” U.S. patent application Ser. No. 08/848,077, filed May 15, 1997 (“the Shear ’077 application”); Collberg et al., “Obfuscation Techniques for Enhancing Software Security,” U.S. patent application Ser. No. 09/095,346, filed Jun. 9, 1998 (“the Collberg application”); and

Shear, “Database Usage Metering and Protection System and Method,” U.S. Pat. No. 4,827,508, issued May 2, 1989 (“the Shear ’508 patent”).

[0052] Broadly speaking, software self-defense is the art of preparing self-defensive programs (SDPs). Because SDPs are intended for execution in hostile environments, they contain various software self-defense measures or mechanisms designed to prevent, detect, and/or report attacks. While software self-defense mechanisms generally do not provide the same level of protection as the physical security of tamper-resistant hardware, software self-defense mechanisms can significantly increase the difficulty of attacks and reduce their probability of success, especially when several software self-defense mechanisms are used together. The appropriate combination of software self-defense mechanisms for a particular application will depend on the nature of the anticipated attacks, the value of the application, the value of the services provided or protected by the application, and/or the hardware environment in which the application is executed.

[0053] A variety of novel software self-defense mechanisms are described below. These mechanisms—such as obfuscation, tamper resistance, and watermarking—typically involve the transformation of an original software program into an “equivalent” software program. Equivalence is not an absolute notion; it has different meanings in different contexts and may be useful for different purposes. In the context of SDPs, however, two object programs are generally considered equivalent if they could have been produced from the same program source code by correct compilers, or if they differ in a manner that would nevertheless be acceptable in compiler output. Accordingly, a transformed program need not be identical to the original program to be considered equivalent, and will generally have some observable differences when executed on a computer. For example, different values may be stored in particular RAM locations, execution times may be different, and/or different results may be computed.

[0054] However, while various aspects of the computer program may be altered by the transformations, if software self-defense mechanisms are to be useful, certain aspects of the computer program must remain substantially invariant. Which aspects of the computer program must not be varied depends on the use to which the computer program will be put. Thus, while strong forms of equivalence are preferred, and indeed necessary in certain instances, in other instances weaker forms of program equivalence will suffice. For example, if a program includes code that receives two inputs, processes them, and generates an output, versions of the program can be watermarked by altering the results of the algorithm. In a first instance, inputs 2 and 2 may yield 4, in a second instance, 2 and 2 may yield 5, and in a third instance, 2 and 2 may yield 6. Under a strict “functional equality” definition of equivalence, and where such equality is necessary, these programs would not be equivalent, as they produce different results for the same input. However, in those instances where one is only interested in whether the result is a positive or a negative number, such programs could be considered equivalent. As yet another example of the relative nature of equivalence, the manner in which numbers are rounded may be significant in some instances, but insignificant in others.

[0055] FIG. 1 illustrates a system 100, such as a workstation or personal computer, for practicing the present invention. As shown in FIG. 1, system 100 might include some or all of the following:

- [0056] a central processing unit (CPU) 104;
 - [0057] a magnetic and/or optical disk drive 106 for reading from and/or writing to diskettes, CDs, or other storage media;
 - [0058] a non-volatile storage unit 108, such as a hard disk drive, for storing data and program files;
 - [0059] a network interface 109 for connecting computer system 100 to a network such as the Internet;
 - [0060] a system memory unit 110, preferably including both high-speed random access memory (RAM) and read only memory (ROM), for storing, e.g., system control programs, data, and application programs loaded from disk drive 106, storage unit 108, and/or network interface 109, such programs including programs incorporating computer code that implements or embodies aspects of the present invention;
 - [0061] a user interface 112, including a display 114 and one or more input devices 116;
 - [0062] a system clock 118;
 - [0063] an optional random number generator 120; and
 - [0064] one or more internal buses 122 for interconnecting the aforementioned elements of the system.
- [0065] The operation of system 100 is controlled primarily by programs stored in system memory 110 (e.g., RAM) and executed by CPU 104. These programs could include:
- [0066] an operating system 146;
 - [0067] a file handling system 148;
 - [0068] a set of user interface procedures 150, for handling input from user interface 112 and for displaying output to the user on display 114;
 - [0069] one or more application programs 152;
 - [0070] a security module 153 for securely transmitting and receiving data to and from network interface 109, display 114, and/or disk drives 106 and 108;
 - [0071] a binary code modification engine 170, described in more detail below, for locating and identifying individual instructions, data, symbol tables, and other elements of a target program 103, and for modifying and/or rearranging these elements;
 - [0072] a software self-defense control program 157, described in more detail below, including without limitation procedures or modules 154 for applying the obfuscation transformations described herein to a target program 103; procedures or modules 156 for adding tamper resistance measures to target program 103; and/or procedures or modules 158 for applying watermarks to target program 103;
 - [0073] software 155 for generating cryptographic keys and/or for performing encryption and decryption, hashing, and/or other cryptographic operations;

[0074] one or more compilers 166 for translating source code files into object code files;

[0075] a debugger 168;

[0076] an optional decompiler (not shown); and a program and/or text editor 174 for creating and/or modifying program or data files.

[0077] It should be appreciated that while FIG. 1 sets forth an exemplary system for practicing embodiments of the present invention, other systems with different architectures and different configurations of subsystems could also be used. For example, embodiments of the present invention can be practiced using only a limited subset of the components shown in FIG. 1. Moreover, while FIG. 1 has been described in the context of a system for performing obfuscation and other self-defensive transformations on a target program 103, it will be appreciated that the architecture shown in FIG. 1 is also generally representative of the type of system on which the obfuscated target program will ultimately run. Accordingly, it should be understood that the system shown in FIG. 1 is intended to be illustrative and not limiting.

[0078] Various software self-defense tools and mechanisms will now be described.

1. Program Obfuscation

[0079] Generally speaking, “obfuscation” refers to a class of software self-defense techniques that produce transformed programs that are more difficult to understand, modify, and/or reuse than the original programs from which the transformed programs were derived. Because attempts to attack or otherwise tamper with a program will generally require an understanding of the program’s functionality, such defensive mechanisms are clearly advantageous. While an obfuscated program may demonstrate different space and time behavior than the original program from which it was derived, the obfuscated program will typically be functionally equivalent to the original program in many respects, and will also typically be expressed in the same format.

[0080] Nearly all software programs can be obfuscated, including, for example, binary program images such as Microsoft® Windows® executables; byte-code images such as Java™ class files; intermediate code forms produced by compilers; assembly language programs; source code representations; and various other program representations. Additional information on obfuscation can be found in the Collberg application, which was previously incorporated by reference herein.

[0081] As mentioned, the job of an attacker is ultimately to understand the function and/or structure of an obfuscated program and to exploit that understanding to his or her advantage. Because few obfuscation transformations or mechanisms are effective against all forms of attack, effective obfuscation often requires a combination of transformations.

[0082] Obfuscating transformations generally fall into one or more of the following categories: (a) those that add complexity, (b) those that remove structure, (c) those that remove useful information, and (d) those that add disinformation. While increasing the space and time costs of a program, a transformation that adds complexity or removes structure will typically make both static and dynamic analy-

sis of the program more difficult, while rendering manual analysis more time-consuming and automated analysis less precise. Similarly, the removal of useful information, such as symbol tables, leaves the attacker with less to work from; while the addition of disinformation, such as misleading symbol tables, can cause an attacker to make incorrect assumptions.

[0083] Although obfuscating transformations can be applied manually as well as automatically in accordance with the principles of the present invention, the automatic application of obfuscating transformations is generally preferable because in most instances such transformations should be applied pervasively and repeatedly to increase their effectiveness. In addition, manual transformations often leave the author with a program that is less understandable and maintainable.

[0084] 1.1. Obfuscation Tools

[0085] A variety of tools and techniques can be used to perform obfuscating transformations on a target computer program. Several exemplary tools and techniques are described below; however, it will be appreciated that any suitable combination of tools or techniques could be used to implement the obfuscation (and other) transformations described herein. Moreover, it will be appreciated that while the tools and techniques presented in this section can be advantageously used to obfuscate a piece of software, these techniques can also be used to implement the other software self-defensive mechanisms described herein.

[0086] As previously noted in connection with FIG. 1, in one embodiment a software self-defense control program is used to perform automatic obfuscation transformations on target programs 103. Software self-defense control program 157 can be formed by linking a library of obfuscation procedures 162 with a binary modification tool 170. The resulting program 157 can thus make advantageous use of the binary modification tool's ability to read and manipulate a target program's binary image in order to perform the transformations specified by library module 162.

[0087] Illustrative examples of well-known binary modification tools include the technologies underlying the Purify memory-use tracing tool, a product of Rational Software Corp. (formerly Pure Atria Software) of 18880 Homestead Road, Cupertino, Calif.; the TracePoint framework for binary instrumentation of programs for test coverage and performance measurement (a product of TracePoint Technology Inc. that is believed to no longer be marketed); the Etch program performance optimization and evaluation system developed at the University of Washington, Box 352350, Seattle, Wash. (see also <http://memsys.cs.washington.edu/memsys/html/etch.html>, and Romer et al., "Instrumentation and Optimization of Win32/Intel Executables Using Etch," Proceedings of the Usenix NT Conference (1997), which is hereby incorporated by reference in its entirety); and EEL, a research project for manipulating SPARC architecture programs conducted by James Larus at the University of Washington (see Larus et al., "EEL: Machine-Independent Executable Editing," Proceedings of the 1995 ACM SIGPLAN Conference on Programming Languages Design and Implementation (June 1995)).

[0088] Binary modification tools typically can—or can with straightforward modification—provide some or all of the following capabilities:

[0089] the ability to locate and identify individual instructions, data, symbol tables, and other program elements, and to communicate this information to the software modules responsible for modifying the program;

[0090] the ability to perform control flow analysis to separate programs from data in non-tagged architectures;

[0091] the ability to identify basic blocks, extended basic blocks, or other units of code execution—a basic block generally being a sequence of instructions having a single point of entrance and a single point of exit, see, e.g., Waite and Goos, *Compiler Construction*, page 329 et seq. (Springer-Verlag, 1984), which is hereby incorporated by reference;

[0092] the ability to identify procedures or other sequences of instructions that are the target of a subroutine or function call;

[0093] the ability to insert instructions before and/or after other identified instructions, and

[0094] to ensure that other instructions are appropriately relocated and updated so that the modified program functions correctly;

[0095] the ability to rewrite and/or modify existing instructions in the program;

[0096] the ability to insert new functions and procedures, including ones of non-trivial complexity (e.g., mapping functions);

[0097] the ability to update symbol information, for example, to regenerate the mapping between symbols and lines of source code in a modified program; and

[0098] the ability to update other program sections, such as the relocation information.

[0099] As described below, the capabilities of a binary modification engine can be advantageously used in concert with other tools to implement software self-defense techniques.

[0100] FIG. 2 provides a more detailed illustration of the functionality and operation of a binary modification tool 200 in accordance with an embodiment of the present invention. As shown in FIG. 2, source code 201 is passed through compiler 202 to yield object code 203. Object code 203 typically includes both instructions and data. Binary code modification tool 200 receives object code 203 and sends code 203 to separator 204, which separates object code 203 into instructions 205 and data 206. Instructions 205 are sent to structure analyzer 207, which separates the instructions into basic blocks 208 and identifies procedures and functions 209. Obfuscation occurs in code modifier 210, which takes instructions 205, data 206, basic block information 208, and procedure/function information 209, and performs a variety of transformations, including, for example, instruction insertion 211, instruction alteration 212, structure alteration 213, and/or any other suitable transformations.

[0101] It is to be understood that FIG. 2 illustrates merely one embodiment of a binary modification tool 200. Binary modification tool 200 may perform additional analyses and/or transformations, and may perform the steps shown in

FIG. 2 in a different order or in an iterative fashion. In addition, while in one embodiment binary modification tool 200 is implemented as a software tool, it should be understood that any suitable implementation of binary modification tool 200 could be used, including without limitation, implementations as a suite of tools, a combination of hardware and/or software processes occurring at different times and/or on different computers, or any other suitable implementation.

[0102] For example, instead of analyzing and operating upon the output of compiler 202, a code modification tool can be integrated within compiler 202. In this embodiment, compiler 202 may include a mode switch allowing compiler 202 to be switched into a mode in which it applies software self-defense techniques to yield obfuscated object code. In another illustrative embodiment, the functionality of a code modification tool is only partially integrated with compiler 202. In this embodiment, compiler 202 is enhanced to generate “hints” for applying program modifications. The hints can act as substitutes for the output of the relatively complex code analysis process performed by binary modification tool 200, so that only the relatively simple modification process needs to be performed by a post-compilation tool. In one embodiment, the hints identify instructions, data, blocks, procedures, and/or functions, thus obviating the need for, e.g., separator 204 and structure analyzer 207. In a somewhat more complex embodiment, compiler 202 includes the ability to generate hints that identify variable locations and references, thus permitting them to be remapped by a post-compilation modification tool—for example, by exchanging locations of two stack variables, then updating all instructions that refer to those variables accordingly. In yet another embodiment, hints are derived from the conventional symbolic debugging information that compilers typically place in binary images (such as the locations of variables, procedures, and the like). Thus, compiler 202 can apply software self-defense techniques directly (e.g., as part of code generation), or can generate information that can be used to implement self-defense techniques subsequent to compilation (e.g., by a linker, a software self-defense tool such as binary program modification tool 200, or other post-processing tools or manual processes). For example, in one embodiment the user of a binary program modification tool 200 could instruct it to perform a particular type of obfuscation, or to apply an obfuscation technique with a particular level of strength, to functions with specified names, or to code that refers to specific variables, or for other specialized purposes; in such a case, tool 200 could be guided by the symbolic name information produced by compiler 202.

[0103] As noted above, compiler integration advantageously obviates the need for some of the complex analysis and decompilation performed by a post-compilation modification tool, since a compiler typically understands where instructions are in the code, how the basic blocks are structured, and so forth. Integration is also advantageous in that a compiler can be more sophisticated about the nature of the software self-defense techniques it employs, and can be used to apply global software self-defense techniques at effectively the same time as more localized software self-defense techniques. For example, a compiler can randomly rearrange the order in which variables are allocated to stack frames or static storage. Such a technique materially

improves the resistance to a comparison of different versions of the program, but can be impractical to implement with a post-compilation tool.

[0104] Compiler integration also has some relative disadvantages. For example, compilation can be expensive. Building a one-megabyte binary image from source code might take several hours, whereas analyzing the source code and applying software self-defense mechanisms might take only a few minutes. Thus, compiler integration can make it more expensive to generate many differently-obfuscated versions of a program. In addition, compilers are often relatively complex. The practical difficulties of adding another function to an already-complex program may, for some applications, merit the use of a separate tool. Compilers are also many and varied. Thus, while integration with a given compiler may be useful for producing obfuscated code of the type produced by that compiler, it may be unable to produce obfuscated code of the type produced by other compilers.

[0105] One of skill in the art will understand that while software self-defense techniques can be advantageously applied to binary or object programs through the use of, e.g., binary modification tools, these techniques can also be applied to programs in symbolic languages, including without limitation assembly language, C, C++, FORTRAN, Java, and Visual Basic, to name just a few examples. For example, to apply software self-defense techniques directly to an assembly language program, an assembler could be created or modified to incorporate software self-defense techniques, or a modification tool could be built that operates on assembly code. For example, binary modification tool 200 could be modified to accept assembly code instead of object code 203, while maintaining the basic functionality shown in FIG. 2.

[0106] If the assembly language is output from a compiler, an assembler capable of creating corresponding machine code can be constructed with relative ease, since such an assembler is typically more straightforward than a general-purpose assembler, as the assembly language code produced by a compiler is typically very constrained (e.g., it typically does not contain macros, header files, or very many symbolic definitions) and more easily processed than programmer-produced assembly language code. Compiler-generated assembly code also typically contains information such as branch labels and function names that enable the relatively straightforward identification of basic block, function, and procedure locations. Thus, as with the forms of partial compiler integration discussed above, the integration of software self-defense techniques into a symbolic language processor, such as an assembler, can eliminate the need to perform complex analysis to locate instructions, blocks, and functions, as this type of information is typically already provided by the compiler and/or symbolic language processor. Moreover, this technique is, to a significant extent, compiler-independent, and thus enables output from different compilers to be processed by the same tool.

[0107] Software self-defense techniques, and especially those that deal with the location of stack data and static data, can also be applied by a source code preprocessor that rearranges the order of data declarations and structure/class members in source code 201. With reference to FIG. 2, such a source code preprocessor may, for example, be applied to

source code **201** before it reaches compiler **202**, or, alternatively, may be integrated with compiler **202**. Although the use of a source code preprocessor can be somewhat expensive (since, e.g., differently-obfuscated versions of a program will need to be compiled separately), it can provide a material benefit by making object code for data references differ from one version of a program to another without having to add valueless instructions. Such preprocessor-based obfuscation techniques will generally work well for structure and class declarations. These techniques also work well for, e.g., conventional C programs, in which all variables in a function are typically declared at the beginning, and can therefore be easily rearranged. Preprocessor-based techniques can be applied independently of build tools, both in terms of how many different versions are created and in terms of which particular source code warrants multiple versions. Moreover, preprocessor-based techniques are typically compiler-independent, and thus do not require access to the compiler's source code. Although preprocessor techniques are somewhat more likely to exercise latent program flaws (since they change the program's references to dynamic data), preprocessor-based techniques have the advantage of generating meaningful symbol tables, thus making it easier to debug different program versions (at least before a post-build tool is applied). Techniques such as self-validation, automated cryptography, obscuring symbol tables by renaming variables, and concealing external calls can also be effectively supported in a preprocessor context, particularly in combination with a post-build tool.

[0108] Thus, it should be appreciated that there are many different ways to implement the obfuscation techniques described herein.

[0109] A variety of obfuscation techniques will now be presented.

[0110] 1.2. Overview of Attack-Specific Techniques

[0111] Obfuscation techniques that are directed at specific types of attacks can be particularly effective in hindering deobfuscation. Illustrative types of attacks, listed in approximately increasing order of difficulty, include:

[0112] 1. Scanning a program's static code and symbol table, looking for strings or instructions that are likely to be important;

[0113] 2. Disassembling a program's code and trying to understand it;

[0114] 3. Running code under a debugger, setting breakpoints, single-stepping, and examining memory and registers.

[0115] The present invention provides systems and methods for countering some or all of these specific threats. For example, as described in detail below, an attack of the first type can be defended against by encrypting the program's code, and/or by eliminating the symbol table information from the binary or replacing it with disinformation.

[0116] The second attack can be defended against by, e.g., ensuring that the code (and the start-up portion in particular) includes hand-crafted segments that result in computed branches, or returns that do not match calls. The code should be self-modifying (e.g., it should overwrite itself as soon as it has been executed) and branch-rich, and preferably should include easily-discovered code that appears to be startup

code but whose primary purpose is to provide disinformation. An array of carefully chosen keys and a loop such as: (a) decrypt region of memory, (b) branch to it, and (c) repeat unless the decrypted region is branched outside, can also be included.

[0117] To combat the third type of attack, code sequences can be crafted that function differently (e.g., branch to a different place) when a breakpoint is set, or when single-stepping is attempted. Because certain debuggers are detectable, either directly or by obtaining information from the operating system, the program can also be configured to determine if a debugger is running, and, if so, to take defensive action. Yet another line of defense is to use the real-time clock to ensure that key instruction sequences are executed in very short periods of time. In so doing, it may be necessary to disable interrupts to keep the program from failing due to normal time-sharing. However, depending upon the circumstances, it is generally acceptable to simply loop until an instruction sequence is executed within a predetermined amount of time, although preferably the loop should not be so tight that it is easy to recognize. Another defense is to periodically compare a location in the code with a location on disk (e.g., in the operating system registry), and to increment both if they agree. In yet other instances, the third type of attack may be defended against by detecting differences in caching or other hardware-dependent behavior exhibited by debugger-controlled code.

[0118] To facilitate a more detailed description of the present invention, novel obfuscation techniques are separated below into the following general categories: local obfuscation, control graph obfuscation, obfuscation through concurrency, data structure obfuscation, and optimization obfuscation. It will be appreciated, however, that any given obfuscation technique may fit into more than one category.

[0119] 1.3. Local Obfuscation

[0120] "Local obfuscation" refers generally to obfuscation that can be safely applied without performing substantial global analysis of the SDP. The following sections describe a variety of local obfuscation techniques.

[0121] Local obfuscation can be performed manually and/or automatically using any suitable code-modification tool. Many of the local obfuscation techniques set forth below are described in terms of modifying the binary object code of a program through the use of a general-purpose binary program modification engine and a set of related processing programs developed for obfuscation. However, one of ordinary skill in the art will appreciate that the present invention does not require use of a general-purpose binary modification engine, and that the described techniques are readily-applicable to code other than binary or object code. It should also be noted that while many of the examples set forth herein are in the context of the Intel® 80x86 instruction set and the Windows® 95, Windows 98, and/or Windows NT operating systems, the techniques set forth herein are readily adaptable to other contexts, instruction sets, and/or operating systems. For example, many of the techniques set forth herein involve control transfers and/or computations for which other architectures have close or direct equivalents. Thus, one of ordinary skill in the art will understand that while their effectiveness may vary somewhat, many of the techniques described in this section and throughout the specification are readily applicable across a broad range of

architectures and/or operating systems, including without limitation Power PC, ARM, MIPS, SPARC, Itanium, the Motorola 68000 family, and M*COR.

[0122] 1.3.1. Local Obfuscation Techniques

[0123] Program customization, micro-obfuscation, and other techniques can be used separately or in combination to construct different versions or instances of a single program and/or to make it difficult to locally understand small portions of code. For purposes of discussion, these techniques can be grouped as follows:

[0124] 1. Techniques that discourage pattern-matching. These techniques focus on constructing different instances of a program that behave equivalently, but that cannot be easily compared or searched for recognizable patterns, thus making it difficult to use knowledge gained from attacking one instance of a program to attack other instances (e.g., by using an automatic tool to “patch” a modified piece of code from one program instance onto the same location in another program instance).

[0125] 2. Techniques that confuse static analysis. These techniques focus on making automated (or even manual) static analysis of the binary code difficult, preferably at both the instruction level and the program level, thus discouraging the use of disassemblers or other program analysis tools.

[0126] 3. Techniques that discourage dynamic analysis. These techniques focus on making dynamic analysis of the code more difficult by confusing what debuggers and/or related tools see at runtime. Some techniques that confuse disassemblers will also fit in this category.

[0127] 4. Techniques that supply disinformation. These techniques focus on supplying “information” that will mislead attackers about the nature of the code.

[0128] 5. Techniques that apply more complex transformations. These techniques typically require more complex analysis than those set forth above.

[0129] Additional aspects of these categories, and the techniques that fall within them, are described below. It should be kept in mind that these categories are being used to organize the discussion of obfuscation techniques. The categories are not rigid, and it should be appreciated that some techniques fit within more than one category.

[0130] 1.3.1.1. Discourage Pattern Matching

[0131] This category of techniques generally involves rearranging and/or substituting instructions or data in the program without changing the program’s function, thus making it difficult to compare two instances of the program and/or to attack one instance of a program using a patch that was successful in attacking another instance of the program. These techniques can be used to supplement many of the control graph obfuscation techniques described below, since debuggers and attackers often assume that related code (e.g., the body of a function) appears close together in memory.

[0132] 1.3.1.1.1. Rearrange Basic Blocks

[0133] This method of discouraging pattern matching involves locating basic blocks in an analysis pass and reordering some or all of these basic blocks pseudo-ran-

domly. The analysis pass may be performed, for example, using structure analyzer 207 of binary modification tool 200.

[0134] FIG. 3 illustrates an application of block rearrangement to a group of basic blocks 310-360 that form part of a program sequence 302, with entry point 300 and exit 370. Program sequence 302 may, for example, be produced by a compiler. Sequence 303 shows these blocks after one possible rearrangement.

[0135] As shown in FIG. 3, each block from sequence 302 is represented by a corresponding block in rearranged sequence 303. It is readily apparent that sequence 303 has the same effect as sequence 302, since the unconditional branches allow the blocks in sequence 303 to be executed in an order that is different from their order in memory. It will be appreciated that such a rearrangement can be performed manually, or by using any suitable tool, including without limitation a binary code modification tool 200.

[0136] A block rearrangement such as that shown in FIG. 3 is typically sufficient to defeat an attempt to compare two versions of the program and/or to use an address-based patch. Block rearrangement is also advantageous when used in conjunction with other obfuscation techniques, since by destroying program locality, block rearrangement makes it more difficult to reverse other transformations.

[0137] 1.3.1.1.2. Split Basic Blocks

[0138] Another technique for discouraging pattern matching is to insert branches into basic blocks, preferably in a random or quasi-random fashion, and to rearrange the resulting block-parts. This insertion and rearrangement may be done transparently by, e.g., the rearrangement function of a binary modification engine, and typically does not require explicit instruction insertion. In one embodiment, conditional and unconditional branches can be used, with instructions added to force the conditional branches to occur.

[0139] An advantage of this technique is that it can be used to split basic blocks into arbitrarily small parts, thus obscuring the code’s flow of control. In addition, splitting long blocks makes it more difficult to locate particular code sequences (sometimes referred to as “code signatures”). However, since signatures can often be quite short (e.g., values of constants and addresses), other techniques can be advantageously employed to further complicate signature detection—such as techniques involving the insertion of null effect instructions or rewriting to obscure operand offsets, as described in more detail below.

[0140] An example of block splitting is shown in FIG. 4, in which control graph 401 represents an unobfuscated control graph and control graph 402 represents the same control graph after splitting. As shown in FIG. 4, control graph 401 includes blocks 400-440, control passing sequentially from one block to another until block 440 is reached. Control graph 401 is transformed into control graph 402 by splitting blocks 410 and 420. Referring to control graph 402, block 410 has been split into blocks 410a and 410b, with an unconditional jump or branch 415 connecting them. Block 420 has been split into blocks 420a and 420b, and a new block 425 has been added, representing an alternative path out of block 420a. Block 420a contains conditional jump 422, with one branch leading to block 420b and the other branch leading to block 425. In order to maintain a control flow equivalent to that of control graph 401, conditional

branch **422** can be designed so that it always branches to block **420b**, and not to block **425** (unless, e.g., block **425** is equivalent to block **420b**). This can be done by, e.g., conditioning branch **422** on an opaque predicate that appears capable of taking on multiple values but actually always resolves to a value or values that trigger the branch to block **420b**. For example, an opaque predicate such as “if x and not x’,” or “if x**2<0,” or any other suitable condition could be used.

[0141] One skilled in the art will appreciate that the obfuscation techniques illustrated in FIG. 4 can be readily used in combination with other software self-defense techniques and methods. For example, the conditional branch from block **420a** can be designed so that it resolves in the direction of block **425** if there is evidence that tampering has occurred or that the program is running under the control of a debugger. In such an embodiment block **425** or subsequent blocks **435** can include code designed to mimic normal operation while taking protective action, as described in more detail below. In another embodiment, block-splitting can be combined with a block rearrangement technique such as that shown in FIG. 3 to provide even greater protection against, e.g., signature-based patching.

[0142] 1.3.1.1.3. Duplicate Basic Blocks

[0143] Yet another technique for discouraging pattern matching is to split a basic block, duplicate the second half of the block, and use a conditional jump to decide whether to branch to the new, duplicate part of the block, or to fall through to the original second half of the block. A related approach is to duplicate entire basic blocks. To enable these transformations, the binary code modification engine should be able to duplicate basic blocks, or to at least adjust their addresses properly if they are duplicated using other means.

[0144] FIG. 5 is an illustration of a block-duplication technique in accordance with an embodiment of the present invention. With reference to FIG. 5, control graph **501** represents an unobfuscated control graph and control graph **502** represents the same control graph after the application of block duplication and splitting transformations **545** and **546**. Block **510** of control graph **501** is duplicated, yielding blocks **510a** and **510b** in control graph **502**. A conditional jump **512** is added to block **500**, and is used to select which of blocks **510a** and **510b** is chosen in a particular execution of control graph **502**. In one embodiment, the condition used in conditional jump **512** includes one or more live program variables. For example, a variable can be compared to a predefined threshold, block **510a** being selected if the value of this variable exceeds the threshold, and block **510b** being selected otherwise. In another embodiment, a non-deterministic condition is used. Both of these approaches make it more likely that conditional jump **512** will advantageously choose different branches on different executions of code sequence **502**. In addition, in one embodiment blocks **510a** and **510b** do not have identical code, so as to advantageously decrease the likelihood that an attacker will be able to recognize that obfuscation is being applied. Proper program execution is ensured by maintaining equivalence between block **510** in control graph **501**, and blocks **510a** and **510b** in control graph **502**.

[0145] Referring once again to FIG. 5, block **520** in control graph **501** is split into blocks **520a**, **520b**, and **520c** in control graph **502**, and block **530** is duplicated to yield

blocks **532** and **534**. In one embodiment blocks **520b** and **520c** are equivalent, but not identical, code sequences. Equivalence between control graphs **501** and **502** is thus maintained regardless of the direction in which branch **522** evaluates. In a preferred embodiment, branch **522** is a conditional branch that includes local program variables that vary between program executions, and/or that depends on a non-deterministic condition.

[0146] In the example shown in FIG. 5, control passes from blocks **520b** and **520c** to blocks **532** and **534**, respectively. In one embodiment, blocks **532** and **534** are equivalent (though preferably not identical) duplicates of block **530** in control graph **501**. In another embodiment, the code contained in blocks **520a**, **520b**, **520c**, **532**, and **534** can be split and/or rearranged in such a manner that the block **520a/520b/532** sequence is equivalent to the block **520a/520c/534** sequence, although neither blocks **532** and **534**, nor blocks **520b** and **520c** are equivalent to each other individually. In yet another embodiment, control passes from block **520b** directly to block **534** via branch **533**, and block **532** is either not created or not used. In other embodiments, other block-splitting and/or duplication arrangements are used. Thus, it should be understood that while the specific arrangements shown in FIG. 5 are provided for purposes of illustration, any suitable implementation of the techniques disclosed herein can be used without departing from the principles of the present invention. Moreover, it should be understood that the block splitting and block duplication techniques described above can be readily applied in combination with other software self-defense techniques.

[0147] 1.3.1.1.4. Rearrange Code Within Basic Blocks

[0148] As mentioned previously, code within basic blocks can be rearranged, thus making it even more difficult for an attacker to identify and/or use patterns and signatures contained within the unobfuscated program. Well-known techniques for local dependency analysis—generally developed to determine allowable instruction re-orderings for use in processor pipeline or multiple functional unit scheduling, see, e.g., Waite and Goos, *Compiler Construction* (Springer-Verlag, 1984), previously incorporated by reference herein—can be used to generate multiple, functionally-equivalent sequences of instructions in a given basic block. These equivalent sequences can be used, for example, to customize copies of the program and/or to serve as watermarks. In addition, such functionally-equivalent sequences can be used to differentiate basic blocks that have been duplicated as described above. Note that for these purposes “functionally equivalent” requires only that the resulting code produce the results expected from the original program, and that it create no undesirable side-effects that alter the overall functionality of the program.

[0149] 1.3.1.1.5. Insert Null-Effect Instructions

[0150] Another technique for discouraging pattern matching is to insert instructions into a program that do not substantively affect the correct operation of the program. This technique is desirable in that it generally does not require knowledge of program semantics to implement. For example, virtually any instruction can be safely preceded by a NOP or null-effect sequence. Moreover, as is the case with code-rearrangement, insertion of null-effect instructions is something that conventional modification tools are capable

of doing correctly. Thus, for example, a program can be customized by quasi-randomly selecting null-effect instructions and/or sequences from a predefined set and inserting them into the program, thereby ensuring that different instances of the program's code will look different. In a preferred embodiment the insertion of null-effect instructions and sequences is used in conjunction with other software self-defense techniques.

[0151] FIG. 6 provides an illustration of the addition of null-effect sequences and instructions to a program. Referring to FIG. 6, pseudo-code sequence 601 moves constants into the CX and BX registers (610a and 610b), moves the value in CX to AX (610c), increments AX (610d), and adds BX to AX, storing the result in AX (610e). Pseudo-code sequence 602 represents program 601 after one possible insertion of null effect instructions and sequences. The techniques shown in FIG. 6 are described below in connection with some or all of the following illustrative null-effect instructions, instruction sequences, and related obfuscation techniques:

[0152] Insert Literal NOPs. Insertion of NOPs 614 before randomly chosen instructions can hinder an attacker's ability to apply patches that work on one instance of a program to other instances, since the NOPs will create different signatures in different instances and cause instructions to have different relative locations.

[0153] Insert Null-Effect Instructions and Sequences. Examples of null effect instructions and sequences include:

[0154] PUSH reg; POP reg (e.g., 616a, 616b in FIG. 6);

[0155] LEA reg,[reg-x]; LEA reg,[reg+x] (e.g., 618a, 618b);

[0156] XCHG reg1,reg2; XCHG reg2,reg1 (e.g., 620a, 620b);

[0157] AND reg, -1

[0158] OR reg, 0

[0159] These instructions and sequences are harder to detect and remove than simple NOPs 614, as it is more difficult to recognize that they do nothing. Note that in choosing and evaluating a potential null-effect instruction or sequence, consideration should be given to the effect, if any, that the instruction or sequence has on flags and special registers.

[0160] Insert Null-Effect Prefix Instructions. This type of instruction has no effect due to the instructions that follow it. In FIG. 6, for example, instruction 612 (i.e., "MOV AX,15") does not affect program execution, since legitimate instruction 610c overwrites the value that instruction 612 moved into the AX register before that value can cause any harm. It will be appreciated that there are numerous examples of other implementations of this technique. For example, without limitation, instructions or sequences such as "LEA reg, value" and "PUSH value; POP reg" could be used instead of, or in addition to, instruction 612. Although any suitable null-effect instructions can be used, two particularly good choices are MOV and LEA, as they

do not affect the flags. Arithmetic instructions can also be used, as long as any significant flags set by the null-effect instruction are also set by the legitimate instruction or instructions that follow. Instructions that involve the AX (or EAX) register or its equivalent are also particularly advantageous because of the various one-byte opcodes for decimal-adjust instructions.

[0161] Replace With Equivalent Instruction. The signature of a program can also be changed by replacing instructions and/or instruction sequences with equivalent instructions or sequences. Examples of this technique include converting shift instructions into rotate and mask instructions, changing constant lengths, and expanding logical operations (e.g., changing "OR AL,0xf0" to "OR AL,0xc0; OR AL, 0x30"),

[0162] Duplicate Idempotent Instructions. Another technique is to replace an instruction with two copies of itself that have the same effect. For example, an AND or OR operation with a constant or variable can be followed by a copy of itself.

[0163] Insert Unusual Instructions. Use of unusual instructions can also be advantageous, as attackers are less likely to take them into account. Examples include DAA, ARPL, SGDT, and flag-setting instructions chosen not to affect correct operation. However, the use of unusual instructions can have undesirable effects on program performance, since unusual instructions often take longer to execute. Accordingly, it will generally be preferable to avoid inserting such instructions in critical loops or other high-performance code sections.

[0164] Insert Null-Effect Suffix Instructions. These instructions have no effect on the operation of the program because of the instructions that they follow. For example, the duplication of idempotent instructions is a special case of this technique.

[0165] Use Computed Offsets. Memory-reference instructions can be rewritten to obscure the actual values of offsets, thus making it more difficult for a pattern-seeking attack to find the correct values. For example, "MOV EAX, offset[reg]" can be replaced with:

[0166] LEA EAX,offset1

[0167] LEA EAX,offset2[EAX]

[0168] MOV EAX,[reg+EAX]

[0169] where offset=offset1+offset2. There are many similar ways to avoid the use of predictable literal offsets or addresses. Examples include performing local constant propagation on the values in registers, and corresponding effective address calculations.

[0170] Rewrite Instructions to Obscure Operand Offsets. When software self-defense techniques are integrated with a compiler, predictable offsets can be avoided by rearranging the order of stack variables or static data before assigning their storage locations.

[0171] Change Base Register. Memory addresses can be developed in multiple steps by, e.g., inserting an LEA instruction prior to the memory reference instruction that changes the base register, rewriting the target instruction to use a different offset, and inserting

another LEA to adjust the base register back. This technique is useful in defeating searches for “interesting” offsets.

[0172] Rewrite to Obscure Constants. Arithmetic and logical operations with constant operands can be rewritten by splitting these operations into two similar operations. For example, one ADD instruction can be split into two instructions, each of which has constant operands that together sum to the original value. The same is true for other operations. For example, a MOV with a constant operand can be replaced with an equivalent sequence involving a MOV with a constant operand and an OR operation with a constant operand.

[0173] Surround Code With XCHG Instructions. Instructions that use registers can be preceded by an XCHG instruction that swaps one or both registers. The original instruction is rewritten to use the different register, and a corresponding XCHG is inserted afterward. This is both a prefix/suffix change, and can surround a larger sequence as long as it is within a basic block that does not use the affected register. This technique advantageously changes the signature of specific register IDs in instructions.

[0174] Adjust Register Values. A target sequence of one or more instructions can be surrounded by a complementary pair of instructions that adjust the value of a register not used in the sequence. This is also a prefix/suffix change that can be used within a basic block. Examples of this technique include the insertion of complementary add/subtract pairs, the insertion of INC/DEC pairs, and the insertion of complementary XOR pairs for any unused constant, stack address, or data address.

[0175] A variety of exemplary null-effect instructions, sequences, and operations have thus been described, any suitable combination of which can be used to obfuscate a program in accordance with the principles of the present invention. Additional examples are described in the Collberg application, which was previously incorporated by reference, and in Cohen, “Operating System Protection Through Program Evolution,” *Computers & Security*, vol. 12, no. 6, pp. 565-584 (Elsevier 1993), which is hereby incorporated by reference.

[0176] 1.3.1.2. Confuse Static Analysis

[0177] The following techniques increase the difficulty of binary code static analysis by operating at two levels. At the instruction level, these techniques confuse disassemblers and other program analysis tools by making it appear that instructions are not what (or where) they should be. At the object file level, these techniques obscure, damage, and/or misdirect high-level structures that are often used to guide automated disassembly and/or other object file manipulation. In so doing, these techniques can frustrate the efforts of attackers using debugging and/or analysis tools such as Sourcerer or SoftICE.

[0178] 1.3.1.2.1. Obfuscating the Executable Format

[0179] Binary editors, and other static analysis tools that an attacker might use, typically must first parse the binary image of a program into instructions and data. In general, this problem is undecidable. Accordingly, the format of the

executable program—e.g., which parts are instructions and which parts are data—lends itself to obfuscation. In particular, programs can be algorithmically transformed so that it is infeasible for a binary editor to precisely parse the transformed binary program image. Initially, the binary editor receives a string of bits with a single address designated as an entry point. The problem is to determine which bits in the binary image are instructions and which are data. The editor proceeds by parsing the instruction at the entry point. To find the next instruction, the editor must disassemble the current instruction and either find its target, if it is a jump (or call, return, etc.), or move to the next instruction if it is not. However, even if an instruction is identified as a jump, finding its target is not always possible (Rice’s theorem). For example, if the target address is computed by a sufficiently complex algorithm, even the application of control-flow and/or data-flow analysis will be unable to determine the jump target address precisely, and will yield an approximate answer instead (e.g., a set of possible target addresses).

[0180] In practice, fairly simple calculations of the target address will suffice to confuse the flow analysis and cause it to yield the set “all addresses” as the possible targets of the jump. For example, the call $f(x)$ can be transformed to $(f+P(\dots))(x)$, where $P(\dots)$ is a function that the obfuscator knows returns 0, but whose computation is sufficiently complex to confuse flow analysis. An even more confusing technique is to choose an arbitrary integer, a , and transform $f(x)$ to $(g+P(\dots))(x)$, where $g=f-a$, and $P(\dots)$ returns a , in order to ensure that the attacker cannot simply guess that P returns 0.

[0181] Other methods of estimating branch targets, such as using relocation information, are also only approximate. Accordingly, transformations such as those set forth above provide provably hard resistance to direct parsing. They achieve similar results to obfuscation in making code difficult to analyze and understand, but they differ from other forms of obfuscation discussed herein, as they rely on instructions and data being indistinguishable without parsing.

[0182] 1.3.1.2.2. Confuse Instruction Recognition

[0183] The following techniques confuse automated and manual instruction recognition, and can be practiced using, e.g., a binary modification engine, such as that described above, which allows manipulation and relocation of branch addresses. These techniques can also be implemented using a two-pass modification tool by inserting markers within the first pass and later changing those markers to the appropriate control transfers.

[0184] Forced-True Conditional Branches. One method of confusing instruction recognition is to insert forced-true conditional branches—either in-line or as replacements for JMPs—and to insert garbage bytes after them that appear to be valid instructions but that actually indicate incorrect instruction boundaries. Forced true conditional branches can be implemented in a wide variety of ways, including, for example: JNZ after OR with a constant other than zero; JNC after any logical operation (e.g., AND, OR, XOR, etc.); and JNO after any logical operation.

[0185] Never-Taken Conditional Branches to the Middle of Instructions. Instruction recognition can also

be frustrated by inserting a conditional branch that will not be taken, the target of which is an inappropriate location. Examples of such locations include the middle of other instructions and locations in data sections. Suitable conditional branches include: JZ after OR with a constant other than zero; JC after any logical operation; and JO after any logical operation. If the branch target is itself an instruction that will not be executed (e.g., because it was inserted for tamper-resistance) or is in a data segment, it can include an immediate constant that appears to be part of a real instruction.

[0186] Use of PUSH, RET instead of JMP. Yet another technique for impeding instruction recognition is to replace JMP instructions with PUSH target-address, RET. This can be done using a binary modification engine. Assuming a free register can be found, the target address can also be generated by a multi-instruction sequence, as described elsewhere herein.

[0187] Obscured Branches. The insertion of branches that target invalid and/or misaligned code can be advantageously combined with the replacement of branches that target good code with relatively obscure equivalent forms that disassemblers cannot easily recognize. This technique makes it difficult for a disassembler to find the target of a legitimate branch, and thus can prevent a disassembler from re-synchronizing itself. For example, a branch through a register that is loaded with a constant, or that is loaded with the result of a computation (e.g., MOV EAX,constant1; XOR EAX,constant2; JMP [EAX];) can be substituted for a simple branch. This modification typically requires analysis of the target instructions to find a register that is not in use and that can be used for the branch target.

[0188] As another example, the following sequence can be inserted after a logical operation:

	JO label1
	MOV EAX,label2
	JMP [EAX]
label1:	garbage instructions
	...
label2:	legitimate instructions

[0189] In this code sequence, the jump to label1 is never taken. However, a disassembler will most likely see the JO and decide that label1 is, in fact, a legitimate instruction address. Garbage instructions can be chosen to make it appear that the location of label2 is in the middle of an instruction, thus causing the disassembler to misinterpret the following legitimate instructions. It will be difficult for the disassembler to resynchronize by analyzing the JMP, since its address is computed.

[0190] RET with Offset. Random bytes can be inserted after a CALL, and the return address of each corresponding RET can be increased to skip over them. This technique will generally make the code look as though the instructions after the CALL will be executed. This technique is particularly effective if the “instructions” after the CALL are fragments

which will cause an instruction-recognizer to become unsynchronized. For example, the following code sequence could be used:

```

CALL Function
Garbage Instruction-1
Garbage Instruction-2
Legitimate Instruction-1
Legitimate Instruction-2
...
Function
...
increment return address to skip garbage instrs
RET

```

[0191] Reference Constants From Other Instructions. References to constants or immediate values in, e.g., MOV and arithmetic/logical instructions can be replaced with appropriately-selected addresses in the middle of null-effect instructions. To implement this technique, a modification engine would insert appropriate constants as null-effect instructions, which then have to be located so that a second pass can update the legitimate instructions that use the constants.

[0192] Replace CALL with JMP to Call-Emulator. A CALL can be replaced with a JMP to one of a collection of locations that emulates the CALL and branches back to the initial JMP after completion. Since JMP does not store a return address, a unique location is needed for each call and return address pair. These emulators can be sprinkled throughout the binary in any suitable location, including as operand bytes in null-effect instructions.

[0193] Replace JMP with CALL and SP Adjustment. If a branch target can be identified reliably, a branch, such as a jump, can be replaced with a CALL by prefixing the branch target with an instruction that pops the return address.

[0194] 1.3.1.2.3. Modify Symbol Tables

[0195] Another set of techniques that can be used to discourage static analysis involves manipulation of the symbol information (and/or the source line map) used by, e.g., debuggers. Symbol tables are generally complex and hard to check, and programs that read them are likely to assume that they are accurate. Thus, as a basic defensive measure, in one embodiment the symbol table is simply removed from the SDP, thus eliminating inter alia the character strings it contains. In another embodiment, the symbol table is damaged by inserting binary garbage (e.g., negative lengths, incorrect offsets, etc.), with the intent of making programs (e.g., debuggers) that attempt to interpret this garbage fail.

[0196] In yet another embodiment, symbol tables are rewritten to be misleading. This is generally a more complex task than damaging the symbol tables, as it is desirable for the results to at least appear valid. Examples of misleading symbol tables include those with randomly-shuffled mappings between variable names and locations, those with variables that have been renamed with valid, meaningful, but incorrect names, and those wherein the mapping of source lines to code has been scrambled. Indeed, if the basic blocks of the program are randomly rearranged in accor-

dance with the techniques described previously herein, the source map will often appear quite random; by replacing it with a source map that has a more conventional, sequential appearance, it will look more normal to an attacker, and thus disguise the application of obfuscation. In one illustrative embodiment, the original source map is used for this purpose—i.e., the map describing the program before the program was obfuscated using e.g., the block splitting, duplication, and rearrangement techniques described above.

[0197] 1.3.1.2.4. Modify Relocation Information

[0198] Generally speaking, the obfuscation techniques described in the previous subsection are also effective with respect to relocation information. Thus, in one embodiment relocation information is simply removed, thereby discouraging sophisticated analysis programs from attempting to analyze the program. Note, however, that this technique is more useful on, e.g., EXE files, as DLL files typically need some relocation information in order to be loaded.

[0199] Similarly, in another embodiment relocation information is damaged. This technique will generally cause analysis programs to fail. Moreover, because the damaged information need not have a relationship to the original information, this technique can be applied to EXE files with relatively little effort. For DLLs, the damaged information preferably includes information which describes relocation of code sequences that are not executed, such as dead code inserted for software self-defense purposes.

[0200] A related technique is to rewrite the export dictionary to include a number of suggestively-named items that are not actually present, thus misleading a human attacker.

[0201] 1.3.1.2.5. Instructions on Stack

[0202] Yet another technique for confusing static analysis involves copying instructions to the stack and executing them there. This technique is a variation of block encryption, which is described in more detail below. An illustrative implementation of this technique is shown in FIG. 7A, in which code sequence 700 is transformed into code sequence 701. Referring to FIG. 7A, an instruction sequence 710 in program 700 is identified that does not include code segment references. In code sequence 701, a jump 720 is inserted at the end of instruction sequence 710 to return to the appropriate location in program 701 (i.e., L2). Code sequence 710 can then be reversed and moved elsewhere in the program, as illustrated by sequence 710a, taking care to adjust any stack references to account for the change in the stack pointer. Next, code 730 is inserted at the old location of instruction sequence 710, code 730 being operable to copy instruction sequence 710 to the stack. A JMP 750 to the stack location of the first instruction to be executed is also inserted.

[0203] FIG. 7B illustrates the execution of code sequence 701. Referring to FIG. 7B, execution of copying code 730 results in instruction sequence 710 and jump 720 being copied to stack 770. Execution of jump 750 transfers control to stack 770, which executes instruction sequence 710, starting with Instr-1. When instruction sequence 710 is finished, jump 720 is executed, returning control back to code sequence 701 at the next instruction to be executed, i.e., POP 4.

[0204] It should be noted that if instruction stream 710 has been encrypted or masked, the copying operation may be

more complex. In such embodiments it may be best to use a CALL to a support routine rather than inserting the copying code in line. Note, too, that while this technique is generally applied only to instruction sequences that do not make code segment references, and is generally limited in its application to individual basic blocks due to branch relocation issues, it is advantageous in that it does not interact with multithreading or multiprocessing. In addition, it will generally confuse static analysis since it is awkward to collect all of the instruction sequences placed on the stack, and it will generally complicate dynamic analysis as well. One skilled in the art will appreciate that this technique is applicable to data segments other than the stack, including the heap or static storage, as long as precautions are taken for multithreading.

[0205] 1.3.1.2.6. Data Allocation Between Code Fragments

[0206] Another technique for discouraging static analysis is to enable dynamic or static data allocation in spaces created between fragments of code. This technique is advantageous, since the existence of modifiable data interspersed with executable code will generally obscure the distinction between code and data, and thus make it difficult for attackers to detect, e.g., self-modifying code. In one embodiment this technique effectively adds parts of the executable program segment into a memory that is treated as an allocable “heap” during the program’s operation. This technique is also advantageous in that it can be used to insert information such as rules, controls, and watermarks into the code to indicate when and where specific actions occurred.

[0207] FIGS. 8A and 8B illustrate an embodiment of a module for creating spaces in executable code and for making these spaces available for allocating data and/or other instructions. The module can also provide operations to allocate and deallocate code spaces dynamically, and these operations can be used to replace selected invocations of the malloc and free commands.

[0208] As shown in FIGS. 8A and 8B, in a preferred embodiment the CodeSpace module includes the following operations:

[0209] Create. In general, spaces can be created within code by, e.g., inserting source-level code or using a binary editor. CodeSpace::create is an inlined operation that creates space statically at the point of call. In the embodiment shown in FIGS. 8A and 8B, the total amount of space created in the code section can be controlled by changing the number of invocations; however, this amount is fixed statically and does not grow dynamically like an actual heap.

[0210] Register. Code spaces created by CodeSpace::create should be registered so that the allocator will be aware of their availability. This can be done in at least two ways: (1) statically, by running a binary editor that looks for the code spaces and puts them in a free list; or (2) dynamically, by having each code space register itself when it is executed. In the embodiment shown in FIG. 8A, CodeSpace::register provides dynamic registration, and is invoked by CodeSpace::create. Thus, a code space created using CodeSpace::create is made available when it is invoked for the first time. Note that for convenience, and unless

otherwise clear from the context, a “code space” will refer to a space allocated using `CodeSpace::create` or an equivalent routine, while a “data space” will refer to a space allocated using `malloc`.

[0211] Allocate. `CodeSpace::alloc` first tries to allocate space from available code spaces. However, if a sufficiently large code space is not available it can allocate space using `malloc`, instead. This feature allows programmers to replace some calls to `malloc` with calls to `CodeSpace::alloc` without having to worry about the availability of code spaces. Since `alloc` may have to call `malloc`, it will generally take more time to execute; however, the performance impact can be reduced by limiting the number of times `alloc` is used instead of `malloc`. Thus, in one embodiment `alloc` is used in place of `malloc` in instances that get invoked relatively infrequently at run time and/or that are likely to be small enough to fit into a code space.

[0212] Deallocate. Spaces allocated using `CodeSpace::alloc` are deallocated using `CodeSpace::dealloc`. In a preferred embodiment, this is true even of spaces which `CodeSpace::alloc` allocated using `malloc`, and thus a programmer need not worry about which kind of space is returned at run time. However, spaces not allocated using `CodeSpace::alloc` should not be deallocated using `CodeSpace::dealloc`.

[0213] FIG. 8C provides an example of a program that uses the `CodeSpace` modules shown in FIGS. 8A and 8B. Referring to FIG. 8C, two blocks of code space are created by `cs.create()` statements 830. Next, the statements `cs.alloc(size)` 832 and `cs.alloc(size+1)` 836 each allocate code spaces for the storage of data, and set the variables `b1` and `b3`, respectively, to point to these spaces. In contrast, the statement `cs.alloc(size+1)` 834 results in a space being allocated using `malloc`, since the amount of space to be allocated—i.e., “size+1”—exceeds the maximum size of the code spaces created using `cs.create()`. Similarly, call 838 to `cs.alloc()` results in space being allocated using `malloc` since the two code spaces created by calls 830 to `cs.create()` are already allocated by statements 832 and 836. Following the call to `cs.dealloc()` 840, however, the code space pointed to by `b1` is deallocated, and thus is available for allocation by the subsequent call 842 to `cs.alloc()`.

[0214] FIGS. 8D and 8E illustrate the operation of the `CodeSpace` module from the perspective of system memory. Referring to FIG. 8D, `CodeSpace`, like other allocators, stores status information 862 in the first byte of an allocated space 860. The caller of `alloc` receives a pointer 864 to the next byte in space 860, thus hiding the status information 862 and protecting it from being overwritten. The part of a space provided to a caller is referred to as the data block 866. In a preferred embodiment, status information 862 distinguishes between at least three kinds of spaces:

[0215] 1. A code space that has not been registered. This is the static state of all code spaces. This information causes `CodeSpace::register` to be called when the create operation is invoked for the first time.

[0216] 2. A code space that has been registered. This information causes the create operation to take no steps on subsequent invocations.

[0217] 3. A data space. This information tells the allocator that it must deallocate this space using the “free”

command, even though it was allocated by a call to `alloc`, since `alloc` called `malloc`.

[0218] In the embodiment shown in FIGS. 8A, 8B, and 8C each code space has the same size. This increases the simplicity of the `CodeSpace` module, however, it will be appreciated that code spaces of different sizes could also be used. For example, better space efficiency can be achieved by managing code spaces of different sizes, and by using a “best fit” policy to select one for allocation.

[0219] Registered but unallocated code spaces are stored in a list that links together the unallocated blocks. In the embodiment shown in FIG. 8D, the list is effectively inlined within the blocks, the first word in each unallocated block storing a pointer 868 to the next such block (or NULL, if there is no such block). Note that unlike status byte 862, the “next” pointer 868 is part of the block provided to callers, which is acceptable because the next pointer is present only in unallocated blocks. Thus, one constraint is that each block be big enough to store a pointer. In other embodiments, such as the one shown in FIG. 8E, free list 870 is a separately-allocated linked list containing (a) pointers 871 to free code spaces 872 and (b) information 874 regarding the size of each such space.

[0220] In a preferred embodiment, the code used to create a space is written in assembly language. This provides close control over the amount of space created, and helps to prevent the compiler from reordering instructions at the boundaries of the created space. Use of assembly code also enables global symbols or labels to be defined in the midst of code segments for purposes of registration.

[0221] The embodiment shown in FIGS. 8A, 8B, and 8C is written in the C++ programming language and should be compiled using the Microsoft compiler, since the source code contains MASM-specific assembler directives. It will be appreciated, however, that other suitable programming languages and/or compilers could be used without departing from the principles of the present invention. Note also that for the embodiment shown in FIGS. 8A, 8B, and 8C the code section in the final binary should be writable. This can be achieved in a variety of ways, including via a source-level run-time function call such as `VirtualProtect()`; or, in a preferred embodiment, via a linker option such as Microsoft linker’s `/SECTION:text,crew` option; and/or via the tool “editbin,” which provides the `/SECTION` option to change existing binaries.

[0222] 1.3.1.3. Discourage Dynamic Analysis

[0223] This category of obfuscation techniques seeks to discourage dynamic or run-time analysis by, e.g., interfering with debuggers’ and/or related tools’ use of breakpoints and tracing mechanisms.

[0224] 1.3.1.3.1. Detection and Obstruction of Debuggers

[0225] To keep an attacker from obtaining useful information about an SDP, it is advantageous to prevent the SDP’s processes from being run under a debugger. For example, periodic checks can be performed to see whether a debugger is being used, and defensive measures can be taken if one is detected. This section describes techniques for detecting debuggers, and for obstructing their operation upon detection.

[0226] 1.3.1.3.1.1. Checking for a Debugger

[0227] A debugger can be detected via operating system interfaces, which typically include calls for checking whether debugger or trace facilities are in use.

[0228] 1.3.1.3.1.2. Performing Environmental Checks

[0229] Another technique for detecting a debugger is to insert calls to a routine (or to one of many quasi-randomly-selected instances of such a routine) that makes environmental checks. For example, a routine can check whether the debug registers are being used to analyze the program and/or whether single-stepping is being performed (e.g., a routine can check hardware-specific registers that cause the micro-processor to trap, single step, etc.). A more sophisticated technique involves the insertion of code to obtain and store the flags, and then to check them elsewhere so as to obscure the relationship between the two operations.

[0230] 1.3.1.3.1.3. Adding Built-in Debugger Process

[0231] Another technique for detecting or resisting a debugger is to split a program into a debugger process and a main process that runs under it. The two processes can communicate via a cryptographic protocol to ensure that they form the equivalent of the original program. Well-known cryptographic techniques can be used to provide mutual authentication. Since most operating systems do not allow a program to be run under two debuggers at the same time, this technique can effectively obstruct an attacker's ability to analyze the program using a debugger.

[0232] 1.3.1.3.1.4. Monitoring Cycle Counter

[0233] Another technique that can be used to detect a debugger is to check the cycle counter to determine whether a predefined piece of code executes within a predetermined number of cycles. In one embodiment, the piece of code is measured in advance—preferably after optimization, customization, and obfuscation—to determine an approximate number or range of clock cycles that should elapse when the piece of code is executed under normal conditions. When run under a debugger, the cycle count will typically be significantly different.

[0234] A related technique is for the program to check the real-time clock to determine whether an appropriate amount of time has elapsed for a measured piece of code. Because real time measurements and cycle counters can be affected by other system activities (e.g., page faults, interrupts, etc.), these indications are statistical in nature and their significance should be evaluated in context. Only when the frequency of measured deviations exceeds a suitably-defined threshold can tampering be reliably inferred.

[0235] 1.3.1.3.1.5. Writing Into Instruction Stream

[0236] Another way to interfere with a debugger is to insert instructions that write data (or other instructions) to a location in the instruction stream that has already been loaded into the processor's pipeline. This technique works best on older systems, such those that use Intel 80386 or 80486 compatible processors, in which instructions are not modified after entering the processor's pipeline. In these systems, an instruction in the pipeline will be executed even if instructions that precede it in the program overwrite or modify the address location at which it was originally stored.

[0237] This technique is illustrated in FIG. 9, in which code sequence 960 includes an instruction 962 which overwrites target instruction 964 with a garbage instruction 968. However, when the program is executed, target instruction 964 will already be loaded into pipeline 966 by the time instruction 962 is executed, and thus target instruction 964 will be executed even though instruction 962 has overwritten the memory location (i.e., 0DA0:010D) at which target instruction 964 was stored.

[0238] Debuggers typically do not replicate this behavior across breakpoints or while single-stepping. Thus, when code sequence 960 is run under the control of a debugger, instruction 962 will overwrite target instruction 964 with garbage instruction 968, and the debugger will subsequently attempt to execute garbage instruction 968 rather than target instruction 964. As a result, the program can be caused to run differently under the control of a debugger, and thus an attacker will have a difficult time obtaining accurate information about the operation of the program. Alternatively, or in addition, garbage instruction 968 can be chosen to signal the presence of a debugger to an external agent (EA), and/or to cause execution to fail.

[0239] 1.3.1.3.1.6. Inserting Race Conditions

[0240] Because debuggers often affect the interleaving order of multi-threaded programs, another obfuscation technique is to intentionally insert race conditions between processes such that these conditions go one way when run "native" but another way when run under a debugger.

[0241] This technique is illustrated in FIG. 10, in which code sequence 1000 contains control threads 1010 and 1012. Threads 1010 and 1012 each execute a sequence of code (1014 and 1016, respectively) before setting a variable, x, to a predetermined value. As shown in FIG. 10, if thread 1010 is able to execute code 1014 faster than thread 1012 is able to execute code 1016, the value of x will be equal to 9 after both threads have finished executing, since thread 1012 will set x to this value, overwriting the value that thread 1010 previously assigned to it. Thus, if a test 1022 is performed after both threads have finished executing (or after both threads have set x), it will be possible to determine which thread set x last. If the order of execution indicates that a debugger is present, appropriate defensive action can be taken (1026). Otherwise, normal execution continues (1024). The order of execution may be different under a debugger, since a debugger will typically execute only one thread at a time, or will stop only one of the threads for single-stepping. It will be appreciated that any suitable technique for monitoring the order of thread execution can be used in accordance with the principles of the present invention. Techniques like that shown in FIG. 10 can make code very difficult to dynamically trace, since if a debugger is detected the behavior of the code may differ from its normal behavior. Such techniques are also advantageous because they can obviate the need for suspicious code fragments that observe the cycle counter or real-time clock.

[0242] Because interrupts and other system processes will occasionally affect thread interleaving, it is typically only possible to give a high probability that race conditions will resolve consistently. However, that assurance can be strengthened by using appropriate loops. For example, block 1026 could test a counter and branch back to the start 1005 unless the counter exceeds a preset amount, in which case an attack could be reported.

[0243] 1.3.1.3.2. Self-Modifying Code

[0244] Another way to discourage static or dynamic analysis of an SDP is to include self-modifying code. Code modification can take a variety of forms and encompasses a variety of techniques, including without limitation patching, encryption/decryption, and interpretation. Code can be modified in place or can be relocated to another location, such as the stack or the heap, and modified there. The following sections present several obfuscation techniques that involve code modification.

[0245] 1.3.1.3.2.1. Dynamically Move Instructions into Stream

[0246] One advantageous obfuscation technique is to insert MOV instructions that store a byte (or some other predetermined amount of data) forward into the instruction stream. The targets of these MOV instructions can be initialized to incorrect values, thus making the static version of the program appear invalid. As the program runs, however, these “errors” will be corrected before they are able to affect execution, since the inserted MOVs will overwrite the incorrect values with the values that should be used during execution.

[0247] FIG. 11 illustrates an implementation of this technique. As shown in FIG. 11, an obfuscated code sequence 1101 is formed by inserting a MOV instruction 1102 into code sequence 1100. MOV instruction 1102 is operable to move an instruction 1103 (or its numeric equivalent) to address location 1104. An instruction 1106—preferably invalid—is stored at address location 1104. When the program is executed, MOV instruction 1102 overwrites invalid instruction 1106 with valid instruction 1103. As a result, when execution reaches address location 1104, it will be occupied by instruction 1103, thus ensuring that execution of obfuscated code sequence 1101 is equivalent to the execution of original code sequence 1100.

[0248] Note that in the embodiment described in the preceding paragraph, only the first execution of instruction 1102 will effectively modify program 1101, since on subsequent executions of instruction 1102, instruction 1103 will already be stored at address 1104. Nevertheless, such frequent stores into the instruction stream can make it more difficult for a debugger to analyze what is happening. Alternatively, an initialization flag can be set after the first execution of instruction 1102. The initialization flag can be tested on subsequent passes, and instruction 1102 can be skipped if the flag is set. Either way, the technique shown in FIG. 11 advantageously lends itself to application in multithreading or multiprocessing environments, as the first thread (or processor) to execute will cause subsequent executions to work correctly.

[0249] The invalid instructions 1106 can be chosen to confuse static analysis. Thus, invalid instructions 1106 may, for example, include the opcodes for instructions of different lengths than the valid instructions that they replace (e.g., INC AX instead of MOV AX, BX), operand specifiers that cause the instruction length of the invalid instruction to be different than the instruction length of the valid instruction (e.g., INC AX instead of INC AH), and/or different operand values (e.g., MOV AX, BX instead of MOV CX, BX). As yet another example, the invalid instructions may replace branches with non-branches and vice versa (e.g., JMP target instead of INC CX).

[0250] To further confound static analysis, an additional MOV instruction 1108 can be inserted in instruction stream 1101 after target instruction 1103 is executed, the additional MOV instruction 1108 being operable to turn memory location 1104 back into an invalid instruction 1110. Such a change ensures that the memory image will be incorrect almost all of the time, thus making the memory image more difficult to analyze. The value 1110 written back to memory location 1104 need not be a constant, but may consist of any available value (e.g., a register value, etc.), since such a variable value will almost certainly not be equal to the correct value 1103. It will be appreciated, however, that the technique described in this paragraph generally does not interact well with multithreaded or multiprocessor operation, since if two threads are executing the same region, one may change the target instruction back to an invalid value before the other thread is able to execute the valid instruction. Thus, when this technique is used in a multithreaded or multiprocessor application, it is preferably applied within a critical section, or to code that is protected by a locking mechanism.

[0251] While modification of the instruction stream is relatively simple within a basic block, application of this technique across basic blocks generally requires that every entry into a block be modified to make the same correction to the block’s code. Application across block boundaries is advantageous, however, in that it involves a JMP or conditional branch before the modified code is executed, further complicating static analysis. In any event, one skilled in the art will appreciate that the transformations of this section are preferably applied after the final basic blocks of the program have been determined.

[0252] 1.3.1.3.2.2. Encryption Techniques

[0253] Another code self-modification technique involves encrypting code sequences and inserting calls to, e.g., a support function that decrypts those code sequences prior to execution. FIG. 12 illustrates an embodiment of this technique. Referring to FIG. 12, portions 1207 of a code sequence 1200 are encrypted to form encrypted code sequences 1202. Any suitable encryption scheme can be used, including, e.g., an XOR with a constant key, a DES encryption scheme, or any of the well-known techniques set forth in Menezes et al., “Handbook of Applied Cryptography,” pp. 191-319 (CRC Press, 1997), which is hereby incorporated by reference. It should be appreciated that entire blocks of instructions (e.g., basic blocks, sets of basic blocks comprising whole functions, etc.) can be encrypted in this manner in the static form of the program, preferably after any other software self-defense changes are made.

[0254] A call 1204 to a support function 1206 is inserted in code sequence 1200. Support function 1206 decrypts code sequences 1202 using a decryption key that reverses the encryption to recover code sequences 1207. The cryptographic key can be stored in a variety of places. For example, the key can be embedded in the call 1204 (at the point of modification), statically embedded in transformation function 1206, or derived from a mapping table in, or accessible to, transformation function 1206.

[0255] Thus, when code sequence 1200 is executed, a call to support function 1206 is made before encrypted code 1202 is executed. The support function obtains the cryptographic key (1205) and uses it to decrypt encrypted code

1202. Decrypted code **1207** is then written into the portion of memory occupied by encrypted code **1202**. Support function **1206** then returns control to code sequence **1200**, and the decrypted code is executed.

[**0256**] The security of this technique can be enhanced by decrypting immediately prior to execution and re-encrypting immediately after execution, thus minimizing the exposure of the decrypted code. Recurring, on-the-fly decryption of critical code also provides more obfuscation than a single decryption at the outset, as it is more difficult to find the decryption code and the keys, and more difficult to determine when and where to look at the result of the decryption. However, in multithreading applications this can require some locking and critical sections within the main program in order to ensure that a thread's decryption and encryption do not interfere with the execution of other threads. However, decryption at first entry also generally requires serializing the decryption function in order to ensure that another entry does not occur during decryption, which, although straightforward, must be explicitly coded.

[**0257**] In yet another exemplary embodiment, the encryption is recursive, with inner decryption routines protected by outer ones. And in one embodiment, the encryption techniques described in this section are used in combination with data allocation techniques similar to those discussed above in connection with FIG. 8, since applying both techniques typically makes the resulting code much harder to recognize than if either technique had been applied in isolation.

[**0258**] FIG. 13 illustrates a variation of the encryption technique described above. Referring to FIG. 13, a portion **1303** of code or data sequence **1302** is replaced with a sequence of invalid data or instructions **1306**. This permits the static form of the program to contain apparently valid (or perhaps confusing) but incorrect instructions.

[**0259**] A one-time pad **1308** is generated by applying a transformation that maps sequence **1306** onto sequence **1303**. For example, the exclusive-or (XOR) function can be used for this purpose, the one-time pad comprising the XOR of the desired instructions **1303** and the invalid instructions **1306**. The one time pad is then stored in memory. In a preferred embodiment, the amount of storage needed to store the one-time pad can be reduced using compression. For example, a Lempel-Ziv encoding, run-length encoding, or any other suitable compression scheme can be used.

[**0260**] Referring once again to FIG. 13, a CALL **1309** to a support function **1310** is inserted in code sequence **1302**. Support function **1310** is operable to decrypt code sequence **1306** using the one time pad **1308**. Thus, support function **1310** will typically include instructions to decompress the one time pad, if it was compressed, and to apply the decompressed one time pad to invalid code sequence **1306**, yielding valid code sequence **1303**. As a result, when code sequence **1302** is executed, support function **1310** causes invalid code sequence **1306** to be replaced with valid sequence **1303**. In one embodiment, a CALL to support function **1310**, or a related support function, is inserted after invalid sequence **1306**, this support function being operable to transform code sequence **1303** back into invalid code sequence **1306**.

[**0261**] Yet another decryption technique is to insert CALLs to interprocess communication functions or MOVs

to flags in shared memory that signal an external process to perform code modification. Specifically, the equivalent of self-modification can be performed using a separate process executing outside the thread requiring the modification. The separate process can be, e.g., another application process, a device driver, or some other independent thread of control. In addition, the separate process can share memory with the modified thread, and poll for flags set in the shared memory region to direct its activities. This technique is advantageous because binary modifications can readily test and set such flags; however, it will be appreciated that any interthread or interprocess communication mechanism available in the system (e.g., RPC, system calls, page faults, etc.) could also be used.

[**0262**] 1.3.1.3.2.3. Prefetch Queue Dependencies

[**0263**] Another code modification technique is to store invalid instructions forward into the instruction stream. For example, HLT or INT 03 may be stored in place of a NOP (or a multi-byte privileged instruction), or a CALL to a support routine may be stored in place of several NOPs. If the cache and prefetch queue already contain the correct instruction, and the processor does not detect data/instruction interactions, the invalid information will not be seen in the instruction stream. However, when single-stepped under a debugger, the stores will complete well before the instructions are actually executed, so the invalid information will appear in the instruction stream. This technique is essentially equivalent to the technique discussed above in connection with FIG. 9.

[**0264**] This technique is generally processor-dependent, and for modern processors, such as the Pentium® processor family, may require that multiple page mappings be available (e.g., set up by program initialization). Once such addressability is established, however, the necessary binary modifications are similar to those for the other instruction stream modification techniques, described above, as are the issues with respect to multithreading and multiprocessing.

[**0265**] This technique may also be subject to random failures caused by interrupts or scheduler activity, since if the processor is interrupted between the store and the execution, execution will fail. For this reason such checks should be considered on a statistical basis, which is why HLT, INT 03, CALL, or a privileged instruction are appropriate choices, since they allow the exception handler (or CALLED function) to note that the exception has occurred, but to resume normal execution unless the frequency of such events is above a predefined threshold. Note that an arbitrary privileged instruction may be preferable to HLT, INT 03, or CALL, since such an instruction would be easily recognizable and would not interfere with other uses of the single-byte exception-generating instructions.

[**0266**] 1.3.1.3.2.4. Implementation

[**0267**] Several of the obfuscation techniques set forth above involve writing or storing to the instruction stream. As a result, these techniques generally need the ability to write to the pages being executed. In some systems that may not be possible with the CS prefix, as some systems do not permit a code segment selector to be created that allows both execute and write permission. However, in Windows 95, Windows 98, and Windows NT, for example, the data segment maps the same addresses as the code segment, but

with different permissions. Thus, obfuscation techniques that involve storing into the instruction stream can be enabled by ensuring that the pages containing the executable image are writeable. As one of ordinary skill in the art will appreciate, depending on system support, that may require adjusting page permissions or copying the image to writeable storage.

[0268] In addition, while some of the techniques set forth above—such as forward stores into the instructions stream and temporary modification of the instruction stream—can be readily implemented as binary modifications to the program, other techniques may require relatively complex code to be included in the SDP, and are thus more easily implemented via function calls. This allows the binary modifications made to the program itself to be relatively straightforward, consisting of little more than calls to the functions that implement the obfuscation techniques, rather than doing the work of the functions themselves. To complicate both static and dynamic analysis of a program that uses such support functions, multiple instances of the support functions can be used, each implemented differently.

[0269] Another implementation issue involves the ability of several of the techniques set forth above to interact with processor-specific handling of writes to the instruction stream. In some Intel 80×86 models (e.g., 80386, 80486), for example, such writes can get “ahead” of the instruction prefetch logic, and may not be recognized for several cycles. In such cases, a JMP instruction can be used to ensure that the prefetch queue is flushed. Moreover, in Pentium and later CPUs, a write to the linear address of a prefetched instruction will cause the prefetch queue to be flushed, so that the modified instruction will be re-fetched when required. Thus, obfuscation techniques should be coded to ensure the appropriate effect with respect to instruction prefetching. In some instances, forcing a branch before the modified instruction is executed will be sufficient. A simple alternative is to follow all such stores with a JMP x+2, where x is the address of the JMP x+2 instruction, which will cause the prefetch queue to be flushed. However, because a Pentium processor’s check of instruction addresses is made by linear address, if two pages of different linear addresses are mapped to the same physical address, it will still be possible to modify an instruction after it has been pre-fetched without the modification having effect. In such cases, a serializing instruction will generally need to be executed, although anything that causes an interrupt or exception will typically suffice. For Intel processors, CPUID is an example of a non-privileged serializing instruction, however, in some applications it may be undesirable to use this instruction since it affects the registers and is not supported on 80486 and earlier processors. However, interrupts or exceptions such as IRET can be used instead, as they effectively serialize as well.

[0270] One skilled in the art will appreciate that although these techniques are described in the context of the Intel 80×86 processor architecture, other processors have similar issues with respect to self-modifying code. For example, RISC architectures generally have less support for code self-modification and typically require serialization or cache flushes in order to have access to modified instructions. However, one of ordinary skill in the art will appreciate that since virtually all processors need some support for code modification in order to support dynamic code generation

and debugging, the techniques set forth herein are generally applicable across a wide range of architectures.

[0271] Obfuscation techniques that involve self-modifying code may also interact with internal and/or external mechanisms that attempt to validate the correctness of the program’s memory image (e.g., the Authenticode® software produced by Microsoft Corporation, or other credential mechanisms, such as those described in the Ginter ’900 patent, the Shear ’754 application, and the Sibert application, previously incorporated by reference herein). While such mechanisms guard against patches, breakpoints, and other modifications, these mechanisms will also typically detect the changes caused by self-modifying code and indicate that they are invalid.

[0272] One way to prevent such an inappropriate detection is to identify the regions in which self-modification may take place, and exempt them from the validity check function. While an attacker could potentially obtain the location of the self-modifying code from the check function itself, this attack can be complicated by encrypting the check function’s validation information. Yet another way to prevent inappropriate detection is to accept multiple results from the check function for certain regions. While this can lead to an increase in validation values (for example, if there are 20 possible self-modifying regions, there could be 2^{20} possible cryptographic hash values for the program, depending on which self-modifications have occurred), this technique is advantageous in that it conceals information regarding the whereabouts of the self-modifying regions.

[0273] A variety of techniques have thus been presented for discouraging static and dynamic program analysis. It should be appreciated that a variety of similar or equivalent techniques could also be used without departing from the principles of the present invention. Because some of these techniques can be relatively expensive at runtime, it may be preferable not to use these techniques in performance-critical code in applications that cannot afford such a degradation of performance.

[0274] 1.3.1.4. Other Transformations

[0275] Additional obfuscation techniques that can be implemented via, e.g., binary code modification are described below.

[0276] 1.3.1.4.1. Transformed External Calls

[0277] In an ordinary compiled program, external calls (e.g., to Windows primitives) can be identified with relative ease: their names through the import dictionary, and their locations through the thusly-identified offsets in CALL instructions. For purposes of obfuscation, it can be helpful to obscure either or both of these relationships. This section describes several illustrative techniques.

[0278] 1.3.1.4.1.1. Funnel External References Through Redirector

[0279] One technique for obscuring external calls is to replace them with calls to a special internal routine that invokes external functions. Referring to FIG. 14, a code sequence 1400 is transformed into obfuscated code sequence 1402 by replacing external CALLs 1404 with PUSH routine-ID instructions 1406 followed by CALLs 1408 to a special handling routine 1410. Special handling routine

1410 is operable to pop the routine-ID off of the stack and call the corresponding external routine in, e.g., library **1412**.

[0280] The routine-ID can comprise any suitable mechanism for identifying the external routine that is to be called. For example, the routine-ID may consist of the name of the external routine or, for greater obfuscation, its address. Because the name or address will not appear with a CALL in the program binary, it can make the job of finding calls to external routines more difficult. Having multiple instances of target code **1410** can make the analysis still more complex, since even if the functionality of one instance of code **1410** is figured out, the other instances will still remain, thus obscuring additional CALLs to external routines in program **1402**.

[0281] A way to achieve even greater obfuscation is to remove all reference to either the name or the address of external routines. For example, in Windows systems the routines in a loaded DLL will typically be in a certain order. Thus, one can examine this order at link time and modify the code of program **1400** to make external calls by ordinal, either directly or by using a redirection technique such as the one shown in FIG. 14. A related technique is to use an initialization routine with an encrypted table of all the names of the routines called externally. The initialization routine decrypts the table, uses run-time calls to locate the relevant entry-points of the external routines, and sets up a table that maps routine-numbers to entry-point addresses.

[0282] 1.3.1.4.2. Automated Decryption of Object State

[0283] Much of the foregoing description has focused on the protection of program code. It is also possible to conceal the object state when it is stored in memory but is not being used. In addition to the techniques described above, one technique for protecting the object state is to use a modification tool to change some or all of the methods of specified classes to decrypt the object state at entry and to encrypt it again at exit or when calls are made out from the method. A relatively simple and fast encryption algorithm such as Vernam can be used, and particular classes can be selected by the programmer for this treatment. Note that the object state need not always be in encrypted form. For example, a boolean flag can be added that tells entry code whether decryption is necessary or not.

[0284] 1.3.1.4.3. String Obfuscation

[0285] It is also desirable to minimize the number of meaningful literal strings in the body of the object program. For example, a program may include strings for user interaction, error messages, trace messages, and the like that could give an attacker useful information about the structure of the program and the location of important routines.

[0286] One way to avoid having meaningful literal strings in the body of the object program is to use a special mechanism such as a “primer” or macro to tag strings so that a modification tool can encrypt them and insert appropriate decryption code where they are used. Again, a relatively simple encryption system like Vernam is sufficient, although other techniques could also be used.

[0287] For example, a macro can be defined as follows:

```
[0288] #define S(str) decrypt_string(SPECIAL_TAG
str)
```

[0289] In this macro the string is prefixed at compile time with the constant SPECIAL_TAG, and passed as a parameter to the decrypt_string function, whose job is to return a pointer to the decrypted string. The programmer (or a special program preprocessor) replaces all uses of constant strings “<str>” in the source code by uses of S(“<str>”). The modification tool searches for the constant SPECIAL_TAG and encrypts the string that follows it. The decrypt_string function preferably maintains a runtime pool of decrypted strings in the heap, and remembers the addresses of the strings so as not to make multiple copies of the same string.

[0290] 1.4. Control Graph Obfuscation

[0291] A relatively strong form of software self-defense can be provided by manipulating a program’s control-flow graph to obfuscate the program’s control flow. There are a variety of techniques for obfuscating a program’s control flow, and these techniques can be performed in a variety of ways, including via manual transformation of the source code and/or automatic source-to-source transformations, obfuscating compilers, or sophisticated binary program editors.

[0292] 1.4.1. Code Replication

[0293] The idea behind code replication is to duplicate pieces of a program’s code such that the structure of the resulting control-flow graph is made more complex, more difficult to generate by a source program using only structured control constructs, and/or one from which it is relatively difficult to recover the original control-flow graph.

[0294] The control-flow graph of the obfuscated program preferably has many redundant paths that correspond to a single path in the original program. Execution may follow any of these paths, as they each produce effectively the same behavior. Thus, different executions of a given path in the original program (e.g., different iterations of a loop) may take different paths in the obfuscated program, complicating attempts at static and/or dynamic analysis.

[0295] The increased complexity of the obfuscated control-flow graph can render the obfuscated program of little help to someone wishing to recover useful source code, as deobfuscating such a program entails discovering which redundant paths have the same behavior. The complexity of solving this problem using static or dynamic analysis makes it difficult to decompile the obfuscated program to a valid source program.

[0296] While this obfuscation technique can be advantageously applied to large, inter-procedural control-flow graphs (e.g., whole programs), it is also applicable to intra-procedural control-flow graphs (e.g., within functions). Indeed, this technique can be applied at the intra-procedural level, and can also be independently and/or simultaneously applied at the inter-procedural level. Thus, while the discussion of code replication that appears below is largely directed to the intra-procedural context, where basic blocks are connected by simple jumps, it will be appreciated that the disclosed techniques are readily applicable to other contexts.

[0297] FIG. 16A illustrates the control graph, G 1602, of a program or program module. Control-flow graph 1602 can be seen as a directed graph with nodes, A_i (1604, 1606, etc.) representing basic blocks, and outgoing edges (arrows) representing transfers of control. Nodes with a single outgoing edge, such as node 1604, represent unconditional jumps, while nodes with two outgoing edges, such as node 1606, represent ordinary conditional jumps, and nodes with more than two outgoing edges, such as node 1612, represent indirect jumps through function pointers or jump tables. Typically, there is a single node with no incoming edges, called the entry node (e.g., node 1603), and a single node with no outgoing edges, called the exit node (e.g., node 1624).

[0298] In most well-formed programs, all of the nodes, A_i , in the input graph are live. That is, all of the nodes in the input graph can be reached by one or more paths from root 1603. In general, each edge is also realizable. In other words, there exists at least one program execution that follows that edge. While an obfuscating transformation need not test for these properties—and indeed, can be applied to programs that do not possess one or both of these properties—it is preferable to substantially preserve these properties to the extent they are true in the original program, as it is desirable to avoid providing clues to a deobfuscator by adding obviously dead nodes or clearly unrealizable edges.

[0299] 1.4.1.1. Code Replication Algorithm

[0300] FIG. 16B illustrates a control graph, H, 1640 resulting from the application of a code replication algorithm such as that set forth in FIGS. 16C and 16D to the control graph, G, of FIG. 16A.

[0301] Referring to FIG. 16C, control graph H is constructed by creating one or more nodes $A_{i,j}$ for each node A_i in control graph G. In a preferred embodiment, a decreasing probability function is used to control the number of nodes $A_{i,j}$ that are added to H, and hence the level of obfuscation (1670-1676). An advantage of probabilistic replication is that it makes it difficult for an attacker to determine how many copies of a node were made. In a preferred embodiment a negative exponential probability function is used, since a negative exponential provides an easy way to bound the expected code growth. In other embodiments, other probability functions can be used.

[0302] Once the nodes of control graph G have been replicated (1678), edges are added to H as shown in FIG. 16D. If a node $A_{i,j}$ in control graph H corresponds to a node A_i in control graph G with a single outgoing edge directed at a target node B_i (i.e., a “No” branch from block 1682), then there will exist one or more target nodes $B_{i,k}$ in control graph H. If there are multiple $B_{i,k}$ in H that correspond to B_i in G (i.e., a “Yes” branch from block 1684) then a test is preferably added to $A_{i,j}$ (block 1686) and edges are connected from $A_{i,j}$ to a suitable number (e.g., 2) of $B_{i,k}$ (block 1690). In a preferred embodiment the particular nodes $B_{i,k}$ to which each node $A_{i,j}$ is connected are chosen quasi-randomly. Otherwise, an edge is simply added from each $A_{i,j}$ to the unique $B_{i,k}$ in H (block 1688).

[0303] The results of this process are illustrated in FIGS. 16A and 16B, in which node 1608 in control graph G 1602 has been replicated by nodes 1644 and 1646 in obfuscated control graph H 1640. Similarly, node 1612 in G has been

copied to nodes 1648, 1650, and 1652 in H. Because nodes 1644 and 1646 have multiple target nodes (i.e., 1648, 1650, and 1652), tests 1645 and 1647 are added to nodes 1644 and 1646, respectively, and edges are connected from each of nodes 1644 and 1646 to a pair of target nodes, as shown.

[0304] Referring back to FIG. 16D, if a node $A_{i,j}$ in control graph H corresponds to a node A_i in control graph G with two outgoing edges directed at target nodes C_i and D_i (i.e., a “Yes” exit from block 1682), there will exist one or more corresponding target nodes $C_{i,k}$ and $D_{i,m}$ in control graph H. The nodes in H can be connected by simply adding edges from each $A_{i,j}$ to a quasi-randomly selected $C_{i,k}$ and $D_{i,m}$ in H (blocks 1692-97). (It will be appreciated that while a preferred embodiment uses quasi-random selection, any suitable method of selection could be used). The results of this process are illustrated in FIGS. 16A and 16B, with reference to nodes 1606, 1608, and 1610 in control graph G, and with reference to nodes 1644, 1646, 1654, 1656, and 1658 in control graph H in particular. A similar procedure can be used to handle nodes with more than two outgoing edges, as is also shown in FIGS. 16A and 16B with reference to node 1612 in control graph G and corresponding nodes in control graph H.

[0305] With regard to exit node, E 1624, in control graph G 1602 there may be more than one corresponding node, E_i 1660, 1662, in new control graph, H 1640. If this is the case, as shown in FIG. 16B, a new node, F, 1664 is preferably created to serve as the unique exit node of H, and an edge is connected from each E_i to F. Alternatively, an arbitrarily-selected E_i , can be selected to serve as the exit node, and the other E_i can be directed to it. It will be appreciated that other ways of handling the exit node could also be used, including simply allowing H (or G) to have multiple exit nodes.

[0306] In selecting nodes $B_{i,k}$, $C_{i,k}$, and $D_{i,m}$, above, care should be taken to avoid creating new dead nodes in the resulting control flow graph. For example, when connecting nodes $A_{i,j}$ to target nodes $B_{i,k}$, it is possible to ensure that there are not any newly-created $B_{i,k}$ without incoming edges by first choosing $B_{i,k}$ that do not yet have an edge from an $A_{i,j}$, if there are such (as long as there are enough $A_{i,j}$). However, this may not ensure that all nodes in the final graph are live, as there can be connected regions of the graph that are not reachable from the root. In FIG. 16B, $A_{1,2}$ is an example of such a dead node, which could be recognized as such by an attacker and discarded. This can be avoided in any of several ways. For example, it is possible to iterate through the nodes of the original control-flow graph 1602 in either depth-first or breadth-first order, starting at the entry node, and deleting nodes and/or adding new edges as appropriate. This will ensure that there is a path from the root to each node $A_{i,j}$, and thus local liveness will imply global liveness.

[0307] When a node A_i has a single outgoing edge to B_i and there are multiple copies of B_i (i.e., $B_{i,k}$), obfuscation can optionally be enhanced by partitioning the group of $B_{i,k}$ into two sets in order to make this case look more like the case where A_i has two outgoing edges. A further refinement would be to use compiler constant-propagation techniques to propagate the results of the tests added to the $A_{i,j}$ through the $B_{i,k}$ to produce inequivalent code segments.

[0308] With regard to the tests that are introduced into the $A_{i,j}$, one embodiment advantageously uses tests that are the

same as, or similar to, tests in nearby code (e.g., tests that are in the same scope, so that the same variables are available). For instance, if there is a nearby test of P , the test $\text{not}(P)$ could be used in an $A_{i,j}$. As another example, if nearby code tests $x > y$, the tests $x < y$, $x < (y+2)$, or $x \geq y$ could be used. Using tests similar to those in nearby code makes it possible for a given test to branch different ways in different executions, thus complicating dynamic analysis and making more edges realizable. Note that these tests or conditions are essentially irrelevant to the correct functioning of the program—it does not matter how they branch, since the nodes or sequences to which they branch are equivalent. In addition, or alternatively, tests and conditions can be chosen that are non-deterministic and/or opaque, in that they are known to the obfuscator, but not the attacker, to always (or never) be true.

[0309] To make recovering a smaller, more useful, control-flow graph difficult, it is desirable to make it relatively difficult to identify nodes in the obfuscated control-flow graph that are copies of a single original node. Thus, it is desirable for the various copies of a node to be relatively different, and/or for copies of different nodes to be similar. To make the copies different, the local obfuscation techniques discussed above can be used, such as choosing different instructions that do the same thing, using different register allocations, inserting short null-effect sequences or opaque tests, and/or inverting the sense of the test instructions at the ends of some blocks.

[0310] In one embodiment the code-replication techniques discussed herein are advantageously applied to relatively large control-flow graphs and relatively small basic blocks. Large control-flow graphs yield more potential execution paths, while small basic blocks make it harder to determine which basic blocks are copies of the same block. It should be appreciated that large basic blocks can be artificially split at a small cost in run time, and/or split differently in different copies of the program or code. This effectively increases the size of the obfuscated control-flow graph without significantly increasing program size, and makes it more difficult for an attacker to find corresponding blocks. It will be appreciated, however, that code-replication techniques are readily applicable to control-flow graphs, and basic blocks, of many different sizes.

[0311] 1.4.1.2. Properties

[0312] The code growth associated with the code replication techniques described above can be adjusted on a per-function basis and/or on a per-block basis, and will typically be bounded by a small multiplicative factor as long as the probability distribution function in 1676 has a small expected value. For example, in the embodiment described above in connection with FIGS. 16A-16D, code growth can be tuned by selecting the probability function used to create nodes in the new control graph, H , and/or by varying the parameters of the probability function that is selected.

[0313] Run-time cost for an interpretive implementation (e.g., the Java Virtual Machine) should be small, as the additional cost is basically the cost of the extra tests. Cache performance, instruction scheduling, and instruction pipeline flushing will be largely unaffected for an interpretive implementation. These issues may, however, become more significant in the context of a native code implementation, but since control-graph duplication is tunable, its impact on frequently-used loops can be minimized.

[0314] 1.4.1.3. Resistance to Attack

[0315] There are a variety of ways to frustrate attacks on a program that has been obfuscated using the code replication techniques described above. First, the number of basic blocks that are added to the control graph can be increased by duplicating blocks, splitting blocks, and/or adding null blocks. In addition, the attacker's task can be made more difficult by enlarging the set of obfuscation transformations that are used, and/or by increasing the strength of the transformations themselves (e.g., by rearranging code, inserting null-effect sequences, and/or making irrelevant conditions suitably stealthy, so that equivalent blocks seem to be different).

[0316] Yet another technique for resisting attack is to split different paths through the control graph in different ways. For example, the block splitting, rearranging, and duplicating techniques that were discussed above can be applied to the control graph. This process is illustrated in FIG. 15, where block X 1530 in the original control graph 1529 is replaced by two equivalent code sections in the obfuscated control graph 1531—one section made up of blocks X_1 and X_2 (1534 and 1536), and the other made up of block X_3 (1538). Block X_1 (1534) includes a conditional branch to X_2 (1536), with a fall-through to block Y_1 (1539) that is not taken. Block X_1 can be written so that it affects the variables that are used by Y_1 (if it branched to Y_1 , which it does not), but not those that are used by X_2 (or at least not permanently and/or adversely). Accordingly, it is only the combination of X_1 and X_2 that is equivalent to X_3 . Analysis can be made more complicated by crossing between paths in such a way that the same block executes in two paths, but plays a different role in each.

[0317] Thus, the code-replication transformation makes it harder to understand the over-all structure of a program or to otherwise reverse-engineer the program, and is able to provide this protection without inserting special security mechanisms into the source code. As a result, this transformation is also valuable for hiding other software self-defense mechanisms.

[0318] 1.4.2. Code Fusion

[0319] Another control graph obfuscation technique is code fusion. The idea underlying code fusion is to find common sequences of instructions that occur in different basic blocks, and to turn these common sequences into "semi-closed" subroutines. That is, the idea is to leave one copy of the common code in place and to replace the other occurrences of the code with branches to the retained copy. Following the retained copy, one or more conditional branches or a jump table can be placed to return control to the code following the other occurrences of the common sequence. Alternatively, all of the common code segments can be copied to a new subroutine which is configured to execute the common code segment and return control to the appropriate locations in the main program.

[0320] Like code replication, code fusion typically introduces irreducible control graphs, thus making the resulting code difficult or impossible to decompile directly into a structured program. Code fusion can also make a program's object code more difficult to understand by making it difficult to recover a control graph that is isomorphic to the original control graph. Moreover, like code replication, code

fusion provides the obfuscator with the flexibility to vary the degree and cost of obfuscation for different parts of the code, and has a relatively small impact on run time.

[0321] Unlike code replication, however, which expands code by creating multiple equivalent instances of basic blocks in the original program, code fusion may contract code—or at least have little influence on its size—because it reuses common sub-sequences of instructions in the original program. However, since jumps or calls are added to break up what was originally straight-line code, there can be a certain amount of performance degradation.

[0322] Code fusion, like code replication, is most easily explained in terms of a program that is represented by a collection of nodes containing basic blocks (e.g., straight-line code not containing branch targets) connected by edges representing branches (conditional or unconditional).

[0323] Known tools, such as those used in the context of optimizing compilers, can be used to perform useful tasks including constant propagation and the computation of node reachability, node dominance (every control path from A ultimately goes through B), and variable liveness. See, e.g., Waite and Goos, *Compiler Construction* (Springer-Verlag, 1984).

[0324] One of the advantages that the obfuscator will have over an attacker is that this analysis can be done in the original, non-obfuscated control graph. Constant propagation and variable liveness are only needed for the extra obfuscation variables, if any, added by code fusion. Thus, it is not necessary to understand or fully analyze a program to obfuscate it using this technique.

[0325] 1.4.2.1. Finding Common Sequences

[0326] Finding replica sets is potentially the most computationally demanding part of the code fusion process. Within the collection of basic blocks, the goal is to find sequences of instructions that occur in two or more places. To that end, it is possible to use virtually any of the well-known pattern-matching algorithms designed to find repeated patterns in large strings, and even the relatively simple ones will work well enough in practice. Since code generated by a compiler tends to be highly stylized, it will generally be fairly rich in common sequences.

[0327] It should be appreciated that for purposes of practicing the present invention it is not necessary to find all common sequences. Instead, it will typically be sufficient to find just enough to obfuscate the structure of the program. In addition, it is not necessary to find maximal sequences, although it is preferable to find sequences at least as long as the branching code that is added. In terms of obfuscation, it will also typically be an advantage to pick sequences that do not correspond to entire statements or expressions, as this will generally make decompilation more difficult.

[0328] 1.4.2.2. Creating Semi-Closed Subroutines

[0329] FIG. 17A illustrates a program fragment before (1701) and after (1702) the application of code fusion. Program fragment 1701 has two basic blocks, A 1704 and B 1712, each of which includes a common sequence of code, X, 1708.

[0330] In transformed program fragment 1702, code sequence X 1708b in block B 1712 has been replaced by a

jump 1720 to code sequence X 1708a in block A. In addition, block A has been modified to include a conditional jump 1722 from the end of code sequence X 1708a back to the appropriate location in block B.

[0331] There are a variety of ways to ensure that test 1722 exits to the correct one of A21710 and B21716. One simple way is to push the appropriate return address (i.e., the address of A2 or B2) at the end of A1 and B1, respectively. In this embodiment, test 1722 simply pops the address off the stack and branches to it. Alternatively, if the number of calls is known—which is the case in the example shown in FIG. 17A—a boolean, integer, or other appropriate variable could be set such that test 1722 branches correctly. A more powerful approach is to use special variables added for the purposes of obfuscation. If an obfuscation variable (OV) has different values in A and B (at least through X), this fact can be used to choose the appropriate branch from test 1722. For example, if OV₇ is 3 in A and 5 in B, test 1722 can compare OV₇ to a constant equal to 4 and use the result to decide which way to branch. As a further example, if some other OV is known to have the same value in both nodes, this fact can also be used by test 1722. For example, if OV₂ is 9 in both A and B, the value of OV₇+5 could be tested against it. Moreover, if no known OV (or combination of OV) has different values in A and B, an OV that is dead in both A and B can be selected and set to different values in A1 and B1, although this may be easier for an attacker to detect.

[0332] In general, it is desirable to use OV that were set a relatively long distance from the test in which they are used (e.g., test 1722). For example, it would be desirable to set the OV to different values in the furthest dominator of A that only reaches B by going through A, and in the furthest dominator of B that does not reach A. Distinct assignments can be made there to the OV—for example, arbitrary small integers can be assigned—and the knowledge of these values can be propagated—not just live to A and B, but available to all nodes reached from A and B in which other assignments have not been inserted. Moreover, if there are other OV known at the point where this OV is to be set, greater obfuscation can be achieved by setting the OV using a value computed from the other known OV, rather than simply assigning it a literal constant.

[0333] Note that an operand is “available” to A from a block B if, in each block in SUCC(B) that leads to A (other than A itself), it is available to A and is not redefined. The “available” state is not generally tracked with live/dead variable analysis, because optimizing compilers are typically not designed to add extra computations, but it can be tracked using essentially similar techniques. In general, it is possible to use either a fresh OV, or an existing OV that is available to A in a node that pre-dominates A and is available to B in a node that pre-dominates B, and is either live in both with distinct known values or is not live in at least one of them (in which case it is set to a distinct value). Using a fresh OV tends to maximize regions of availability, while reusing an OV presumably makes it harder for an attacker to keep track of what the OV really records. Moreover, to minimize overhead (and clues), if there is an OV that is live or available in either A or B, but dead in the other, a distinct value can be assigned to it using a single assignment as far back as possible on the dead side.

[0334] The code fusion technique described above does not get much more complicated when the common sequence

appears in three or more basic blocks. However, when the branching factor gets large enough, it is advantageous to convert the branching controls from a sequence of conditional branches to a jump table. To increase stealth, this should be done when the branching factor reaches approximately the size at which a compiler would convert to a jump table.

[0335] The values of OV's that are used in conditional branches are relatively unconstrained, and will be stealthier if they are not all taken from a small set. However, when a jump table is used, it is generally advantageous to have a reasonably compact set of possible values for a single OV. The generation of these values can be simplified by processing the largest replica sets first, and assigning contiguous or nearly contiguous values to their OV, thus enabling the construction of relatively compact jump tables. In addition, a few non-reachable values can be placed in such a jump table to provide disinformation.

[0336] Note that a degenerate case of code fusion occurs when X is empty. In this case, code is not actually overlaid or fused, but complexity is nonetheless introduced into the control flow graph, as more basic blocks, joins, and conditional branches will be present.

[0337] Although it has been implicitly assumed that A and B can be broken up without worrying about, or affecting, their meaning, this may not be a safe assumption when the condition code (e.g., part of the hardware state) is live between instructions. Accordingly, some local analysis should be performed—using, for example, conventional defined-used bookkeeping—to ensure that instructions are not inserted that change the condition code at a point where it is live (i.e., used before being set again).

[0338] It has also been implicitly assumed that there is no overlap between the replicas of different replica sets. It should be understood, however, that it is possible to allow overlap, although this typically complicates the task of the obfuscator without necessarily imposing a correspondingly greater burden on the attacker. In addition, while the above description has treated X as straight-line code, it should be understood that this need not be the case. It will be recognized, however, that if X is not straight-line code, it will generally be necessary to make each exit branch conditional on the path of entry.

[0339] 1.4.2.3. Relation to Code Replication

[0340] Code fusion can be readily applied in combination with code replication. However, while there is no reason in principle why, for example, code fusion cannot be applied to the output of code replication, code fusion might well eliminate many of the redundant nodes introduced by code replication, thus lessening code replication's effectiveness. In addition, while the larger and more complex control graphs resulting from code replication will typically raise the cost of the code fusion transformation, code fusion should not significantly raise the cost of a subsequent code replication. Accordingly, it will generally be more advantageous to apply code fusion first, followed by code replication.

[0341] In one such combination of code fusion and code replication, the code replication process is modified to only generate multiple instantiations of nodes that are the targets of unconditional branches, leaving nodes that are the targets

of conditional branches alone. That is, conditionals are only added to the end of basic blocks that do not already end in a conditional. This will be referred to as Code Replication Lite.

[0342] FIG. 17B shows another variation of code fusion that can be used to further confuse a potential attacker by “crossing branches.” As shown in FIG. 17B, this involves swapping block B21736 with block A21734, so that code sequence 1731 contains A1, X, and B2, while code sequence 1733 contains B1, X, A2.

[0343] Since most compilers eliminate branches to unconditional branches, and unconditional branches to branches, it is advantageous to do the same with the branches introduced by code fusion, as this will make it more difficult for an attacker to identify them.

[0344] 1.4.2.4. Additional Protective Techniques

[0345] There are a number of ways to strengthen the protective techniques described above. For example, one general technique is to make it hard to identify the obfuscation variables. This can be accomplished in a variety of ways, including:

[0346] Avoid Allocating the OV All Together. For example, a macro can be used to allocate various pieces of storage for OV, and to mark them with a primer, a la watermarking.

[0347] Avoid Always Using OV in Exit Tests. For example, regular program variables can be used to make the necessary distinctions, with constant propagation (perhaps only local) informing the obfuscator which variables/values will suffice.

[0348] Another protective technique is to follow code fusion with Code Replication Lite. The obfuscator replicates the successors of unconditional exits from basic blocks, and uses arbitrary tests on OV in the conditional branches. If the OV that is tested is not constant in an entry node, then making the exit branch unconditional will change the set of realizable paths in the control graph. Another protective measure is to implement the exit test of one or more of the replica sets as a jump table containing some holes that are filled with the addresses of arbitrary basic blocks.

[0349] Yet another way to strengthen code fusion is to replace some unconditional branches with opaque tests on OV, either instead of, or prior to, applying Code Replication Lite. The non-taken branch in these opaque tests can go to an arbitrary basic block, although it is preferably a block with existing, reachable code that is relatively distant from the test in the control graph.

[0350] Yet another illustrative strengthening technique is to do full constant propagation and live/dead analysis on a set of the ordinary program variables, instead of using variables from a separate OV set. For the exit test of each replica set, a variable that is known to have distinct values (or to be dead) at all entries can be used, adding a fresh variable only if such an ordinary variable cannot be found. This technique can entail a significant amount of work for the obfuscator, however, as it is generally a much more delicate task than propagating a special set of obfuscation variables or constants and keeping track of which of these variables are live, available, or dead. In short, the obfuscator will have to understand the program to obfuscate it using this

technique. However, this technique increases the difficulty of attack by making virtually every conditional branch in the code a candidate for replication, and by making it unsafe for the attacker to prune an exit the first time, or any other time, control reaches a breakpoint. In addition, such a technique can obviate the need to follow code fusion with code replication.

[0351] 1.4.3. Partial Code Fusion

[0352] Code fusion has been described in terms of its application to identical code sequences; however, this is only one limiting case. Code fusion can also be applied to code sequences that are not identical, and indeed, can be applied to code sequences that are totally different (another limiting case).

[0353] FIG. 18 illustrates the application of code fusion to code 1802. Code 1802 includes several code sequences 1804, 1806, 1808 that are similar to each other, but differ in specific locations. For example, code sequence 1806 includes the same sequence of instructions as code sequence 1804, but some of the operands are different. Similarly, code sequence 1808 is similar to a portion of code sequences 1804 and 1806.

[0354] Code sequence 1810 illustrates one possible application of partial code fusion to code sequence 1802. As seen in FIG. 18, code sequences 1806 and 1808 have been replaced with calls (1816, 1818) to a semi-closed subroutine 1811. In the example shown in FIG. 18, before each branch (e.g., 1816, 1818) to subroutine 1811, the locations that differ between code sequence 1811, on the one hand, and 1804, 1806, or 1808, on the other, are patched with the values appropriate to that branch. This can be done, for example, by using a sequence of MOV instructions with immediate operands, as illustrated by code sequence 1812. An alternative that may be preferable if there are a relatively large number of differences between the code sequences is to use a loop that copies values from a patch table. This technique is illustrated by code sequence 1814, which includes a loop that moves values stored in patch table 1817 to locations pointed to by address pointers that are also stored in the patch table. It will be appreciated that other patching techniques could be used without departing from the principles of the present invention. As with other self-modifying code techniques, partial code fusion generally requires a region of memory that is mapped as both code and data.

[0355] When partial code fusion is applied to a program—in the manner illustrated in FIG. 18, for example—the patched locations can be effectively treated as obfuscation variables, and techniques similar to those used for other obfuscation variables can be used to decide where to set the patch values. For example, if it is practical, it can be advantageous to place the assignments that set the patched locations relatively far from the associated branches in the control graph. This is illustrated in FIG. 18, in which intervening code 1815 separates patch code 1814 from branch 1816.

[0356] Moreover, when there are more than two calls to a subroutine, as shown in FIG. 18, not all the values may need to be patched for each call. The necessary assignments can be determined by constant propagation in the same, or substantially the same, manner used for other obfuscation

variables. Thus, in FIG. 18, code sequence 1808 can be replaced by code sequence 1812, where the operands y3 and y4 are not patched into subroutine 1811, those operands having been previously patched when code sequence 1814 was executed. Of course, if code sequence 1810 were designed to change subroutine 1811 back to its original form before executing code sequence 1812, then code 1812 could include additional MOV instructions to patch y3 and y4 into the appropriate locations of code 1811 (i.e., the locations of arg1 and arg2, respectively). One of ordinary skill in the art will appreciate that other variations could be used without departing from the principles of the present invention.

[0357] As shown in FIG. 18, subroutine 1811 includes a branch_test 1822 that is operable to return execution to the appropriate location in the program. The location to return to at the end of the subroutine can be selected in a variety of ways, including without limitation:

[0358] Testing a conventional boolean or integer variable.

[0359] Testing one of the patched values. For example, branch_test 1822 can test the argument of instr-5. If the argument is z3, then branch_test 1822 jumps to the next valid instruction following code sequence 1812 (i.e., instr-y). Similarly, if the argument of instr-5 is z2, branch_test 1822 jumps to the next valid instruction following code 1816, and so forth. Moreover, if the program contains additional subroutines similar to subroutine 1811, obfuscation can be increased by using patch values from one or more of those subroutines in branch_test 1822. Constant propagation can be used to select these patch values such that branch_test 1822 branches to the proper location.

[0360] Patching a return address directly into the branch instruction. For example, patch table 1817 can contain the address of the appropriate location following code 1816, such that code sequence 1814 is operable to write the correct return address directly into the appropriate location of branch_test 1822. In such an embodiment, branch_test 1822 could consist of a simple jump instruction, the destination of the jump instruction (i.e., RESUME-a) being patched in by code sequence 1814.

[0361] Patching the return branch itself. An illustration of this technique appears in code sequence 1812, in which instruction 1824 is operable to patch a jump instruction directly into subroutine 1811. Thus, when subroutine 1811 is called by instruction 1818, an instruction sequence equivalent to sequence 1808 will be executed, and control will jump from code 1811 to instr-y.

[0362] The patched locations can be initialized in a variety of ways. For example, if there is one call to the subroutine that is dominated by the program's entry node (without control passing through the other calls), the patch locations can be statically initialized to its values. Thus, in FIG. 18, the locations in code sequence 1811 that will be subsequently modified by code sequences 1814 and 1812, respectively (e.g., the locations of x1, x2, instr-3, etc.), can be initialized to the values used by code sequence 1804, since that is the first of the relevant code sequences that is executed.

[0363] Alternatively, or in addition, branch addresses can be initialized to locations that are in data, or in the middle

of instructions, thus misdirecting attempts to perform static analysis of the program. The correct branch addresses can be patched in by the calls to the subroutine, as described above. Similarly, other operand addresses can be initialized to code locations to confuse code/data discrimination, and/or op-codes can be initialized misleadingly, for example, by replacing branches with data operations and vice-versa.

[0364] Moreover, if the branch that terminates the subroutine is patched into the subroutine during execution, analysis can be complicated further by initializing the location to be patched (and one or more following locations, if desired) to contain arbitrary and/or misleading code. This will obscure the end of the subroutine, as it will typically be difficult to discover that this code is not executed. Thus, in FIG. 18, for example, a portion of branch_test 1822 can be initialized to contain confusing values (assuming, of course, that branch_test 1822 contains code to pass execution of original code sequence 1804 on to the next appropriate code sequence).

[0365] A variety of code sequences lend themselves to partial code fusion. For example, in-lined functions and standard code sequences for common control structures (e.g., loop headers, function calls and returns, etc.) are typically well-suited for application of this technique. Another good candidate for partial code fusion is table-driven self-checking code.

[0366] Indeed, code sequences that perform table-driven patching are themselves good candidates for replacement by semi-closed subroutines. This is illustrated in FIG. 19, which shows a code sequence 1900 to which partial code fusion has been applied in the manner set forth above. Code sequence 1900 includes code sequences 1910, 1912, 1914, and 1916, each of which makes use of a patch table to write appropriate values to the subroutine to which it branches. Code sequence 1902 is formed by applying partial code fusion to code sequence 1900. In the example shown in FIG. 19, code sequences 1910, 1914, and 1916 are replaced with calls to code sequence 1922 and with appropriate patches. For example, code sequence 1916 is replaced with code sequence 1926, which branches to code sequence 1922 after patching the number of items to be patched (i.e., the loop limit, max) (1930), the address of the patch table (e.g., a table similar to table 1817 in FIG. 18) (1932), and the starting address of the appropriate subroutine (1934). Thus, it should be appreciated that partial code fusion can be recursively applied an arbitrary number of times.

[0367] A variety of additional techniques the same or similar to those set forth above can be used to make partial code fusion more obscure. For example:

[0368] If there are live obfuscation variables at the point where a patch is performed, the values of those variables, as determined by constant propagation, can be used to generate the patching values, rather than implementing the patching values as immediate operands or directly stored values.

[0369] The patch tables can be encrypted, to be decrypted upon use. This can be accomplished in a variety of ways, such as those set forth previously. For example, simple XOR keys can be generated as functions of obfuscation variables, or if two blocks of code, A and B, are available, their XOR can be used to convert either block to the other.

[0370] Code fusion generally complicates self-checking less than many other forms of self-modifying code. If constant propagation is used to keep track of the patched values, it is possible to ensure at obfuscation time that any given self-check skips patch locations whose values are not fixed at the point of call of that particular self-check, and it is also possible to calculate the effect on the hash function of locations whose values are fixed. Moreover, the semi-closed routines can be collected in one area of memory if it is desired to bound the range of locations where self-modification occurs, for purposes of self-checking, for example.

[0371] 1.4.4. Dynamic Code Overlay

[0372] Prior to the advent of virtual memory, programs that were larger than available physical memory were executed using overlays. In accordance with this technique, the logical address space of the program was partitioned into a collection of sets of overlays. The members of a set of overlays were linked to execute in the same physical address space, and hence only one member of a set of overlays would be physically resident at any given time. When control was to transfer to an overlay that was not currently loaded, the target overlay was read from disk and was used to replace whatever overlay was currently loaded in the address space to which it was linked.

[0373] FIG. 20 illustrates the use of overlays in accordance with one embodiment of the present invention. Code and/or data segments 2006, 2008, and 2010 are stored in memory 2002. For example, segments 2006, 2008, 2010 could represent sequential sections of a program. In a preferred embodiment segments 2006, 2008, and 2010 each contain code to load the next overlay. For example, code can be included at the end of segment 2006 that is operable to cause segment 2008 to be loaded into memory. Program 2004 is stored elsewhere in memory. When program 2004 is executed, overlays 2006, 2008, and 2010 are loaded into memory region 2012 as appropriate. When an overlay is loaded into memory region 2012, it overwrites any overlay that was previously loaded at location 2012. Thus, through the use of overlays, a program can be designed such that its entire executable image does not appear in contiguous memory locations. Also, because different segments of the program will appear at the same memory location, the use of overlays can make address tracing more difficult. Additional information on overlays can be found in Levine, *Linkers and Loaders*, chapter 8 (Morgan Kaufmann 1999), which is hereby incorporated by reference.

[0374] Overlays were extensively studied in the 1960s and 1970s, and many algorithms for efficiently partitioning code into overlays were published. See, e.g., Knuth, *The Art of Computer Programming*, vol. 1, pp. 453 et seq. (1968). For present day purposes of obfuscation, however, it is not necessary to be as concerned with efficiency, since overlays can be read in (and optionally decrypted) from some other part of main memory, rather than from disk, and because a high degree of overlay is not needed to conserve main memory as was necessary in the past. For example, in FIG. 20, overlays 2006, 2008, and/or 2010 can be stored in main system memory (e.g., RAM) along with executable image 2004, rather than being stored separately on disk. Thus, it will be appreciated that relatively simple overlay algorithms will suffice for purposes of obfuscation, including those that

can be implemented with a binary code modification engine. For example, the techniques described above for patching locations in semi-closed routines can readily be adapted to copy in overlays.

[0375] A technique that is related to the use of overlays is to simply move pieces of code occasionally to freshly allocated space, and to either shuffle some other piece of code into the old space, or to place the old space on a free space list. As illustrated in FIG. 21, code can be inserted into a program that is operable to allocate a new memory space (2102); copy a section of code to the newly allocated memory space (2104); patch entry and exit branches from the old memory space (i.e., the memory segment from which the code was copied) to the new space (2106); and deallocate the old memory space (2108). Thus, when execution of the program reaches the point where it would have branched to the old memory space, it will branch instead to the new memory space, execute the code contained there, then branch back to the appropriate memory location of the next part of the program (e.g., the location to which the code that was previously-stored in the old memory space would have branched). For purposes of allocating and deallocating space, the techniques described above in connection with FIG. 8 can be used. It should also be appreciated that moving code generally requires patching up relocation entries in other code that accesses the moved code by absolute address. Accordingly, in a preferred embodiment the relocation table and its use are hidden or obscured using the techniques set forth in the Collberg application or elsewhere herein.

[0376] 1.4.5. Protecting Jumps

[0377] If an attacker identifies a critical piece of code that tests some condition and takes different actions based on that condition, he can focus his attack on that test. For example, in FIG. 22A test 2202 might check to see if a password is correct. Conditional jump 2204 (e.g., JZ, JNZ, etc.) branches to block A 2206 if the test evaluates one way (e.g., password correct), and branches to block B 2208 if the test evaluates the other way (e.g., password incorrect). For example, block B 2208 could be a sequence of code that prints out an error message such as "Password Incorrect" and exits the program, while block A 2206 could be the code sequence that is executed if the password is correct.

[0378] A technique for strengthening this test against attack is shown in FIG. 22B. A jump 2212 corresponding to jump 2204 in FIG. 22A is repeated for each of its possible outcomes. Thus, if jump 2212 evaluates one way, execution is routed to another jump 2214 which performs the same, or substantially the same, evaluation as jump 2212. If jump 2212 evaluates the other way, then execution is routed to jump 2216, which also performs the same, or substantially the same, evaluation as jump 2212. If the two consecutive jumps do not evaluate the same logically, then the jumps preferably branch to a tamper response mechanism 2218. As a result, the attacker has to defeat at least two jumps. If only one jump modified, the tamper response code gets executed.

[0379] Note that obfuscating code 2213, 2215 can be added between jump 2212 and jumps 2214 and 2216, respectively, in order to make the relationship between those jumps less clear. Such code should, however, preserve the condition code for the second jump, either by avoiding

instructions that set it, by saving and restoring it as necessary, and/or by making the second condition code a predictable function of the first.

[0380] It will be appreciated that tamper response code 2218 can be implemented in any suitable manner, including, for example, using the techniques set forth herein. If the tamper response is made strong enough to disable the code permanently, an attacker that simply tries a few things is likely to get caught by this mechanism.

[0381] An extension of the technique shown in FIG. 22B is to chain N protected jump blocks together, as shown in FIG. 22C where N=3 has been selected. The depth of this construction, N, is the number of jumps that would have to be modified to defeat this mechanism. Moreover, it should be understood that while in one embodiment each of the jumps tests the same or substantially the same condition code, this need not be the case. For example, in one embodiment, jumps 2224 and 2226 may execute normally (i.e., reach block A) if a particular condition is true (e.g., JZ), whereas jumps 2225 and 2227 may execute normally (i.e., reach block B) if a related condition is true (e.g., JNZ). In other embodiments, this relationship can be varied. For example, jump 2224 and 2226 may be inverses (e.g., JZ and JNZ, respectively), with normal execution occurring if the condition code of jump 2224 is true and the condition code of jump 2226 is false (e.g., a "Yes" and a "No" exit from jumps 2224 and 2226, respectively).

[0382] Another variation that provides additional resistance to attack is to intertwine separate applications of these techniques. For example, arrangements such as those shown in FIGS. 22B and 22C can be combined to share one or more common levels of jumps, as shown in FIG. 22D, in which the second level of jumps is shared (i.e., jumps 2238 and 2240). This can be supported by a PUSH/POP combination, as shown, in which code to pass control back to the appropriate location of the program (e.g., code 2234 and 2236) is effectively saved (e.g., pushed onto the stack), to be executed after flow has passed through the final level of jumps. In one embodiment, shared jumps 2238 and 2240 can test for either of two conditions. For example, jump 2238, or an equivalent code sequence, may be operable to jump to block 2242 if either condition A or condition B is true. In one preferred embodiment, tests A and B (2230 and 2232) are operable to produce the same condition code (e.g., set the same flag or variable), and jumps 2231, 2233, 2238, and 2240 are identical.

[0383] In other embodiments, more than just the conditional jump is repeated. For example, one might reevaluate and test expression 2202 in FIG. 22A using an independent code sequence. Another variation is to perform a checksum or hash on the original code that is responsible for evaluating the condition, in order to make sure that it has not been tampered with.

[0384] 1.5. Obfuscation Through Concurrency

[0385] Concurrent (multi-threaded) programs are often difficult to design, debug, and understand. This property can be exploited for the purposes of obfuscation and tamper resistance. The basic idea is to hide a deterministic program in a non-deterministic state machine, where components of the state are modified by multiple threads of control in such a way that the correctness of the program is preserved by

construction. An attacker who does not know the invariants of the state, and the partitioning of state components into obfuscation variables and significant variables, faces a potentially exponential number of possible interleavings (or genuine concurrency) in the various threads.

[0386] 1.5.1. Baton Passing

[0387] 1.5.1.1. Basic Baton Passing

[0388] This section describes a technique that allows an obfuscator to develop a program with the ease associated with the development of sequential programs, while presenting an attacker who tries to understand the program with the complexity associated with relatively-unstructured, concurrent programs.

[0389] Given a program composed of a series of sequential statements (e.g., $S=S1; S2; \dots Sn$), a fresh variable, PC, can be added to the program, and the series of statements can be rewritten more obscurely as:

```

PC = 1;
while PC <= n do {
  if PC == 1 then {S1; PC = 2}
  if PC == 2 then {S2; PC = 3}
  ...
  if PC == n then {Sn; PC = n+1}
}
```

[0390] The loop in the code sequence shown above executes until PC reaches a predetermined value, n+1. The if-statements contained in the loop are operable to execute one of the statements, S_i , from the original sequence, and to increment PC. Thus, PC is effectively “passed” from one if-statement to another until all of the relevant statements S_i in the original code sequence have been executed.

[0391] For more obscurity, the if-statements can be permuted arbitrarily:

```

PC = 1;
while PC <= n do {
  if PC == 5 then {S5; PC = 6}
  if PC == n then {Sn; PC = n+1}
  ...
  if PC == 2 then {S2; PC = 3}
}
```

[0392] And to make the code even more difficult to follow (e.g., with a debugger), several threads can be started at the same time, and the relevant statements S_i can be partitioned among them. For example:

```

PC = 1;
parallel {
  /* thread 1 */
  while PC <= n do {
    if PC == 5 then {S5; PC = 6}
  }
  |
  /* thread 2 */
  while PC <= n do {
```

-continued

```

    if PC == n then {Sn; PC = n+1}
  }
  |
  ...
  |
  /* thread m */
  while PC <= n do {
    if PC == 3 then {S3; PC = 4}
    if PC == 2 then {S2; PC = 3}
  }
}
```

[0393] In the example code sequence shown above, each thread contains a code sequence similar to the code sequences that were discussed previously. Since PC is shared by all of the threads, it can be used to ensure that the sequence of statements S_i is executed in the proper order. Incrementing PC in one thread can pass control to another thread (or simply pass control to another if-statement within the same thread). The obfuscator can still be confident in the correctness of this parallel program, based on an analysis of the original (simpler) sequential program, and based on the fact that PC is a fresh variable, not modified by the S_i . In effect, PC implements a baton passed among the threads, such that the multiple threads emulate a single virtual thread, VT, in which the statements execute in their original order.

[0394] This technique generalizes to programs that are not simple sequences of statements. For example, branches can simply be treated as assignments to PC that do something other than increment it, while conditional branches become conditional assignments to PC.

[0395] To make the original structure even more difficult to understand and/or reconstruct, in one embodiment the program state is augmented with an assortment of obfuscation variables, OV, and obfuscation constants, OC. These can be of any type used in the program, although it will generally be helpful if OV is disjoint from the “genuine” program variables, $GV+\{PC\}$. Furthermore, an assortment of obfuscation statements, $OS_1, \dots OS_k$, can be created. In one embodiment obfuscation statements are made to appear relatively unobtrusive (i.e., to appear much like genuine statements, S_i). In general, obfuscation statements can read any available variables ($GV+\{PC\}+OV$) and constants ($GC+OC$); however, in a preferred embodiment they should write only to obfuscation variables (OV), terminate (e.g., no loops or recursive procedure calls), and avoid causing run-time errors or infinite loops.

[0396] Obfuscation statements can be scattered in each of the threads. For example:

```

Pc = 1;
parallel {
  /* thread 1 */
  while PC <= n do {
    if PC == 5 then { S5; PC = 6}
    if PC == 2 then { OS1; OS5}
  }
  |
  /* thread 2 */
  while PC <= n do {
    if PC == n then { OS2; Sn; OS7; PC = n+1; OS3 }
```


-continued

```

    }
    |
    ...
    |
    /* thread m */
    while PC <= n do {
        if PC == 1 then { OS1; OS5 }
        if PC == 3 then { S3; PC = 4 }
        if PC == 5 then { OS2; OS5 }
        if PC == 2 then { S2; PC = 3 }
        if PC == 3 then { OS3; OS6 }
    }
}

```

[0397] Thus it is possible to place obfuscation statements within segments relatively freely, and also to add new top-level segments containing obfuscation statements. As far as the virtual thread is concerned, these are just null-effect sequences since they terminate cleanly and do not adversely affect genuine program variables. But to an attacker they can create non-trivial, concurrent computational activity that can make it difficult to distinguish the virtual thread's activity from obfuscation activity.

[0398] 1.5.1.2. Distributed Baton

[0399] In the scheme described above, a baton is passed to a segment as an atomic action—for example, by a single assignment to a primitive data type like integer or boolean. As discussed in more detail below, however, it is not necessary that cancellation of a segment's possession be atomic, since the segment does that itself; however, the cancellation should generally be complete before the hand-off occurs.

[0400] If a Gray code is used, setting a single bit would be enough to pass the baton, but $\log_2(n)$ bits (plus some dummies) would have to be tested. However, this generally holds true only if S_1, \dots, S_n are to be executed in sequence. If there is a more elaborate control structure, more bits might have to be changed.

[0401] A related technique is to replace PC with a collection of “guard” variables, g , one for each of the original segments. In one embodiment boolean variables are used, and to increase obfuscation a similar collection of obfuscation guard variables, og , are also used in obfuscation segments. The guard bits are preferably manipulated so that no two genuine segments have the baton (e.g., no two guard bits are true) at the same time, and so that there is always one genuine segment that either has the baton or is already being executed. Passing the baton is the last genuine code in a genuine segment. For example:

```

g1 = true; g2 = false; ...; gn = false; g0 = false;
og0 = false; og1 = true; ... ogn = false;
parallel {
    /* thread 1 */
    while !g0 do {
        if g4 then { g4 = false; S4; g5 = true }
        if og4 then { og4 = false; OS1; OS4; og5 = true }
        if g6 then { g6 = false; OS6; g7 = true }
    }
}

```

-continued

```

/* thread 2 */
...
|
/* thread m */
while !g0 do {
    if gk then { OS2; Sn; OS7; gk = false; OS19;
                g0 = true; OS3 }
}
|
/* thread m+1 */
while !g0 do {
    if og1 then { og1 = false; OS1; OS4; og4 = true }
    if g3 then { S3; g3 = false; g4 = true }
    if og4 then { og4 = false; OS2; OS4; og5 = true }
    if g2 then { og3 = false; g2 = false; S2; g3 = true }
    if og3 then { og3 = false; OS3; OS6; og4 = true }
}
}

```

[0402] In this example, control is passed between (and within) multiple threads by making assignments to guard variables, g_i . For example, a conditional statement (e.g., an if-statement) in the last thread shown passes execution of the virtual thread to a conditional statement in the first thread by setting g_4 to true. When g_4 is true, the conditional statement in the first thread evaluates to true, and a sequence of instructions is performed which executes a statement from the virtual thread (e.g., S_4) and updates the guard variables appropriately. Obfuscation statements, obfuscation variables, and/or obfuscation threads can be added to obscure what is going on.

[0403] To add additional obfuscation, a compile-time key can be used to determine whether “selected” corresponds to true or to false for each g_i . That is, the key determines whether the guard-bit tests are for g_i or $!g_i$, the assignments of true and false being reversed depending on the i th bit of the key. (Or, if the g_i are not boolean variables, the key can determine, for example, how g_i should be assigned so that a test involving g_i evaluates to true or to false).

[0404] In the example set forth above, genuine branches can safely set obfuscation guards, and obfuscation branches can safely read genuine guards. Moreover, in this example, each branch assigns to two guard bits. It will be appreciated, however, that other conventions could be used to further complicate an attack. For example, each branch could write to, e.g., at least three guards. In the case of an obfuscation branch, all three guards could be obfuscation guards, and for a genuine branch the guards could comprise two genuine guards plus one obfuscation guard.

[0405] An attacker's analysis can be complicated still further by selecting a subset of the obfuscation guards to be invariant (although assigned to), and using them freely in the tests, knowing the outcome. In a genuine branch they can be (re-)assigned their invariant values. While in an obfuscation branch whose guard is (or includes) an invariantly false obfuscation guard or the negation of an invariantly true one, it is possible to not only assign complementary values, but also to assign to genuine guards (indeed, to genuine program variables), making it harder to distinguish the sets. In addition, arbitrary extra guards can be conjoined to the tests for obfuscation branches, and tautologous (e.g., invariantly true, or the negation of invariantly false) extra guards can be conjoined to the tests for genuine branches.

[0406] 1.5.1.3. Deeper Invariants

[0407] One way to increase obfuscation is to distribute the baton over a collection of variables, and to encode the baton using a set of invariants that makes it hard to draw inferences about the uses of particular variables, but that still preserves the property that the baton can be passed easily with a single atomic action. Two exemplary techniques are described below, namely the use of obfuscation invariant guards and the use of virtual PCs.

[0408] 1.5.1.3.1. Obfuscation Invariant Guards

[0409] As previously indicated, it is generally desirable to obscure the relationship between genuine guards, on the one hand, and obfuscation guards and obfuscation invariant guards on the other. This can be done, for example, by intermixing these variables in a thread's tests, and/or by varying the sense of the variables, as discussed above. To further increase obfuscation, it is desirable to be able to assign to a genuine guard variable in an executable obfuscation statement, thus making it more difficult for an attacker to determine which statements are obfuscation statements, which statements are genuine statements in the virtual thread, and which guard bits are genuine.

[0410] The present invention provides an advantageous technique whereby an obfuscator can assign incorrect values to genuine guard variables without causing incorrect execution of the virtual thread. Specifically, error correcting codes can be used to correct "erroneous" guard variable values. Thus, genuine variables can be assigned erroneous values (from the perspective of proper execution of the virtual thread), the errors being removed through the application of error-correcting codes before they are able to derail proper execution of the virtual thread. It should be appreciated that any suitable error-correcting code(s) could be used for this purpose, including without limitation Hamming codes, Reed-Solomon codes, and the like.

[0411] 1.5.1.3.2. Virtual PCs

[0412] Another technique for distributing the baton is to use virtual PCs. For example, a collection of k integer variables can be used to represent PC. The table of k-tuples that is actually used to represent the virtual PC (VPC) values of each virtual thread (VT) is preferably an obfuscation-time secret. That is, the obfuscator can make it up, use it to obfuscate the program, and throw it away. It need not appear in the program.

[0413] For example, a table for k==3 might be:

VPC	VPC1	VPC2	VPC3	sum product	
1	2	3	4	9	24
2	2	5	4	11	40
3	2	6	4	12	48
4	2	6	7	15	84
5	2	6	8	16	96
6	3	6	8	17	144
7	4	6	8	18	192
8	4	9	8	21	288

[0414] It will be appreciated that while, in this example, the table has been kept fairly regular for ease of explanation, the table can be readily constructed in a different manner.

For example, in one embodiment some randomness or quasi-randomness is used in the construction of the table.

[0415] Since, in this example, each of the VPC_i is a positive and monotonically non-decreasing function of VPC, and one of them increases at each step, the sums and products of the triples are monotonically increasing, and therefore distinct. Moreover, since only one of the VPC_i changes at each step, VPC can be incremented by assigning to a single VPC_i . (Note, however, that in general an arbitrary transfer of control requires using a mutually exclusive or mutex lock around as many as k assignments.) An illustrative application of this technique is shown below, using the relationships set forth in the VPC table shown above:

```

VPC1 = 2; ... VPC2 = 3; ... VPC3 = 4 ...;
OV9 = 31; OV11 = 75; OV13 = 9;
...
parallel {
/* thread 1 */
while (VPC1+VPC2+VPC3) <= OV9 do {
if (VPC1+VPC2+VPC3) == 15 then {S4; OS2; VPC3 = 8}
if (VPC3*OV13) == 18 then {OS1; OS5; VPC1=1}
}
}
/* thread 2 */
...
/* thread n */
while (VPC1<=4) do {
if (VPC1+OV13+VPC3)==19) then {OS2; S5; VPC1=VPC2/2}
}
}
/* thread n+1 */
while (VPC2*VPC3)<OV11 do {
if VPC2 < 4 then { OS1; S1; OS5; VPC2 = VPC3+1}
if (41<(VPC1*VPC2*VPC3)<63) then {S3; VPC3=3*VPC1+1}
if VPC1 == 2 then { OS2; OS5 }
if VPC2 == 5 then { S2; VPC2 = VPC2 + 1; }
if (VPC1*VPC2*VPC3)==36 then {OS3; OS6; VPC2=VPC3-1}
}
}
}

```

[0416] It should be appreciated that the variable names (e.g., VP1, VP2, VP3, OV9, OV11, etc.) used in this example (and in the other examples shown herein) have been chosen to aid the explanation of this technique. In a preferred embodiment these variables are given obscure and/or misleading names in order to obscure their functionality and/or interrelationship. For example, the names of some or all of the variables could be arbitrary strings, or could suggest a relationship with an unrelated or irrelevant operation or code segment.

[0417] It should also be noted that while, in the example shown above, the baton is passed through the VPC table linearly from the lowest VPC to the highest (i.e., the first genuine statement, S1, corresponds to VPC=1, the second statement, S2, corresponds to VPC=2, etc.), this need not be the case. Instead, the table can be stepped through in an arbitrary order, and in the general case any suitably pre-defined virtual thread could be used. For example, assignments to the VPC_i could be made in a manner such that the first statement, S1, executed if VPC was equal to 1; S2 executed if VPC was equal to 3; S3 executed if VPC was equal to 2; and so forth. In this case, the assignments to VPC_i would be made so that the value of VPC varied sequentially from 1, to 3, and then to 2.

[0418] As shown above, the values of VPC can be tested without clearly revealing the values of the individual VPC_i to static analysis (e.g., by testing the sum and/or the product of the VPC_i). Moreover, it is possible to test for a range of values rather than for an exact value (e.g., $41 < (VPC_1 * VPC_2 * VPC_3) < 63$). This can be especially advantageous when testing product-terms. In another embodiment, the VPC itself can correspond to a range of values, thus allowing some assignments to be made to the VPC_i without changing the value of VPC. Since the sum (or the product) can be tested, and since the value of VPC can be changed by an assignment to, e.g., a single VPC_i , an attacker will generally have difficulty distinguishing between genuine segments and obfuscation segments.

[0419] Some rows in the VPC table can be uniquely determined by the value of a single VPC_i . In the example shown above, $VPC == 1$ and the evenly-numbered rows are so determined. This can provide yet another way to test the value of the VPC without revealing all that is known about the VPC_i . For example, following the test $VPC_2 == 3$, a genuine segment can increment VPC by the assignment $VPC_2 = 5$, and an obfuscation segment can leave it unchanged by the assignment $VPC_3 = 4$.

[0420] Moreover, code that is never to be executed (and which can therefore contain arbitrary code) can be guarded by any test that is not satisfied by any of the rows. Without knowledge of the table, static analysis generally cannot reveal the unreachability of this code. For example, referring to the table shown above, it will generally be difficult to statically determine that, e.g., the tests $(VPC_3 * OV_{13}) == 18$, $(VPC_1 * VPC_2 * VPC_3) == 36$, or $VPC_3 == 5$ are not satisfied. In addition, one or more versions of the code that establishes some or all of these impossible conditions can be scattered in other never-executed segments, so that in order to establish that any one of them is not executed, an attacker generally must know that all of them are not executed.

[0421] It should also be noted that the VPC_i need not be statically initialized to values corresponding to a value of VPC. For example, each of the VPC_i could be initialized to, e.g., -1, and in three different real threads' main loops, conditional assignments similar to the following could be placed:

[0422] if $VPC_1 == -1$ then $VPC_1 = 2$;

[0423] if $VPC_2 == -1$ then $VPC_2 = 3$; and/or

[0424] if $VPC_3 == -1$ then $VPC_3 = 4$

[0425] Moreover, these tests and assignments need not be simple or deterministic. For example, initialization of the VPC_i could be made contingent on a test of whether the low order 3 bits of the clock are zero, a test of whether a predefined loop is executed in less than 100 microseconds, a test of whether a predefined real thread has repeated its main loop at least 150 times, a test of whether a pointer returned by malloc is a multiple of 32, or any other suitable test that is repeatedly executed until it succeeds (and preferably is expected to eventually succeed within a predetermined period). If something suspicious is detected, on the other hand, the value of VPC_i can be set to, e.g., 0 or 1, thus preventing proper initialization and effectively preventing the virtual thread from starting.

[0426] It should be appreciated that the segment associated with a VPC value need not be contiguous in a real

thread's main loop. Since there can be multiple ways to test for a single value of VPC, one such test can be performed, part of the associated segment can be executed, and VPC can be left unchanged. Later, in the main loop, the same VPC value can be tested for in a different way, and execution of the segment can be resumed. This process can be repeated, VPC being set to its next value after completing the final piece of the segment. As another example, if the top-level control construct of a segment is a fixed loop on an integer variable (e.g., i from 1 to 10), incrementing the loop variable can be done unconditionally in the main loop of the containing real thread. The segment detects the first time through (by the variable being over limit, e.g., 11) and initializes the variable for the current execution (e.g., to 1). It increments VPC only after the execution with the limit value (e.g., execution with $i=10$).

[0427] Thus, the above-described baton-passing techniques generally make static analysis more difficult, due in part to the use of implicit control transfers and due to the fact that few tools are available to perform multi-threaded data-flow analysis. These techniques can also be effective in impeding dynamic analysis, due to the presence of obfuscation segments and due to the difficulty involved in debugging multiple threads. These techniques can be implemented at the source-code level as well as at the binary or object code level.

[0428] Potential disadvantages of baton-passing techniques are that the use of multiple threads tends to slow program execution, and the resultant code can be difficult for the obfuscator to debug. In addition, data structures used by different segments of the virtual thread generally cannot be allocated on a stack, but instead must be allocated statically or on the heap so that all threads can access them.

[0429] However, these potential disadvantages can be substantially avoided or reduced by limiting the use of baton-passing techniques to certain well-chosen places. For example, self-checks are typically particularly good candidates for the application of baton-passing. In general, self-checking code includes code for checking for debuggers and also checking the integrity of other pieces of the program. Baton-passing can advantageously be applied to self-checking code, whereby segments that perform some crucial work are intermixed with segments that make self-checks and segments comprised of executable and/or non-executable obfuscation statements, the intermixing making it more difficult to remove the self-checks without interfering with the crucial work. Accordingly, if it were determined that for a particular application a relatively widespread application of baton-passing obfuscation techniques would lead to unacceptable performance degradation, baton-passing obfuscation could be applied on a more limited scale to those regions of code that perform self-checks, or any other suitably chosen regions, thereby obtaining the advantages of baton-passing while lessening the effect of its relative disadvantages.

[0430] 1.5.2. Race Conditions

[0431] Multi-threaded programs are generally constructed to avoid race conditions that result in non-deterministic behavior. However, as previously described, it is possible to build race conditions into a program that resolve in the correct order under normal conditions, but not when the program is executed, e.g., under a debugger. Moreover, race

conditions need not resolve in the correct order all of the time; instead, code can be put in a loop and repeated until it succeeds or until it becomes sufficiently probable that the repeated failures are due to deterministic behavior.

[0432] Another use of race conditions is to create a relatively long loop in the virtual thread of baton passing described above. The loop is preferably made to appear as if it has no exit, while, in fact, it exits if an improbable (but not impossible) race condition occurs. This technique can advantageously make static and/or dynamic analysis more difficult. An argument against the use of such techniques is that when constructing a difficult-to-solve maze, it is generally desirable to construct the maze such that there is only a single, relatively long path through it. The application of non-determinacy can create multiple paths, and, unless care is taken, an attacker may by chance find a relatively short one. Careful design is necessary to ensure that even the “luckiest” paths are still adequately long.

[0433] 1.5.3. Non-Determinism

[0434] Debugging non-deterministic programs—even those without race conditions—is generally harder than debugging deterministic programs, since replaying program execution is typically more difficult. As a result, non-determinism can be advantageously injected into a program to render its operation more obscure to an attacker. Referring to FIG. 23, for example, in one embodiment branches such as branch 2302 are inserted into the program, the branches resolving in a substantially non-deterministic fashion. The program statements to which the branches lead (e.g., blocks 2304 and 2306) preferably consist of obfuscation statements and/or equivalent (but preferably non-identical) genuine code sequences. Thus, in FIG. 23, blocks 2304 and 2306 can be designed such that it does not matter which block is chosen by indeterminate branch 2302, since the program will execute correctly either way.

[0435] One source of pseudo-non-determinism is the cycle counter in Pentium and more advanced 80x86 processors. The cycle counter is typically reset when the processor is reset—e.g., at power-up—and ticks with each clock cycle. At 233 MHz, for example, the low-order 32 bits of the cycle counter wrap around in approximately 18 seconds, and the low-order 16 bits wrap around in about 0.3 milliseconds. In practice, it will be difficult for an attacker to eliminate this randomness, as enough randomness is typically present in the low-order 16 bits due to differences in disk read times (typically 10 ms), of which there are usually many when booting the operating system, and due to differences in the actual times when clock interrupts happen (typically every 50 ms). As a result, the value contained in the low-order bits on a given reading of the cycle counter can be treated as effectively random. Reading the cycle counter typically takes only one instruction; however, since such an instruction is usually easy to recognize, in one embodiment it is guarded by self-checking code and/or generated dynamically, e.g., on the stack.

[0436] 1.5.4. Resisting Attack from a Known Point in the SDP

[0437] One potential point of attack on an SDP is its startup code. A determined attacker will be able to figure out where the operating system transfers control when starting up an executable, and will also be able to determine the state

of memory at that point. With knowledge of the pre-startup condition, the attacker can simply start the program and attempt to follow along from there.

[0438] As a result, it is desirable to lose an attacker before going near a sensitive part of the SDP. That is, it is desirable to make it difficult for the attacker to follow execution of the program. Concurrency can be of assistance in this regard. For example, a sensitive segment of the SDP can be given a prologue whose primary or sole purpose is to elude and/or detect attackers. In a preferred embodiment, such a prologue uses a substantially different collection of defensive techniques than those used by the sensitive code itself, so that if an attacker is able to crack parts of the prologue, he will gain at most minimal clues about the defensive techniques used in the sensitive code.

[0439] FIG. 24 illustrates one technique for interfering with an attacker’s ability to determine what is happening in a sensitive piece of code. As shown in FIG. 24, several real threads (2410, 2412, 2414) can be started, which taken together implement a concurrency-obfuscated virtual thread (PVT) 2402. Some or all of these threads can primarily execute statements 2416 which are inserted for obfuscation and have no apparent purpose, or which misleadingly suggest that they serve some genuine purpose. Since a typical attacker will not know in advance that some or all of the threads are playing obfuscatory roles, he will have to assume that each thread is performing a legitimate function and trace it. As the number of threads—and the complexity of each thread—increases, the attacker’s task will typically be rendered more difficult.

[0440] To prevent an attacker from simply removing threads, one or more genuine statements and/or segments can be embedded in some or all of the threads. Thus, if the attacker removes a thread, the program will not execute properly and/or other defensive action will be taken. Although the genuine segments may do little real work, in a preferred embodiment they should at least perform busy work or otherwise emulate obfuscation segments.

[0441] In one embodiment the virtual threads (2410, 2412, 2414, etc.) have different sets of batons and disjoint sets of genuine variables, yet share obfuscation guards and obfuscation variables. Each virtual thread preferably loops a relatively large number of times without using an obvious loop counter. For example, the loop can use a compiled-in pseudo-random sequence generator to generate a new value on each iteration, and the loop can do something different when the resulting value equals the nth value in that sequence (as computed in advance by the obfuscator). The loop can then perform part of the virtual join, discussed below, and continue looping so that it is not obvious that anything of special significance has happened.

[0442] The PVT 2402 terminates (and virtual join 2418 occurs) when each of the component threads (2410, 2412, 2414) “finishes” execution (e.g., when the VPC for each thread reaches its predetermined value). When this occurs, the exit virtual thread (EVT) 2403 is established (2420, 2422) and executed. The guards of EVT 2403 are preferably set one at a time by statements in the component threads of PVT 2402 (for example, after each thread has executed a loop a predetermined number of times). In one embodiment an EVT guard is set to a distinctive value by PVT 2402 when too much real time has elapsed since the prologue was

started (e.g., a predefined value such as a few hundred milliseconds, or any other suitable value for the given application). Thus, if any of the threads, $O_1 \dots O_N$, (2410, 2412, 2414) is delayed by an attacker, the prologue will never get around to exiting, as the guard value can, e.g., preclude virtual join 2418 and/or 2430 from occurring.

[0443] In one embodiment genuine statements in the component threads 2410, 2412, and 2414 of PVT 2402 are operable to store words that become all or part of the code of EVT 2403. In addition, genuine statements in PVT 2402 are also preferably operable to construct a key 2419 that can be used to decrypt the main body of EVT 2403. Each real thread, $O_1 \dots O_N$ (2410, 2412, 2414) and $E_1 \dots E_N$, (2424, 2426, 2428) can also keep a count of the number of times it has gone around its loop, and from time-to-time compare this with the count(s) maintained by some of the other real threads. If any real thread gets unreasonably ahead or behind (e.g., by more than can be accounted for by vagaries of scheduling), another thread can exit its loop, thus spoiling the other threads' chance to complete. Note that since a debugger typically executes one real thread at a time (or all but one), each virtual thread should have one or more segments in each real thread.

[0444] In a preferred embodiment a second phase of the attempt to lose an attacker involves the use of stronger, or different, forms of cryptography. EVT 2403 can use keys (e.g., 2433) constructed by its threads (2424, 2426, 2428) to decrypt components (2432) of code. The code can be concurrency-obfuscated, although it is operable to play more than an obfuscatory role. In addition, virtual threads may decrypt pieces of code for their own segments that will be executed later, and overwrite code for their own segments that will not be executed again, as well as decrypting other keys for future use.

[0445] After the trail has been sufficiently muddled, various virtual threads can start performing non-obfuscatory work. For example, they can look at the log to see how many unsuccessful—and not obliterated—attempts there have been to execute the sensitive code since the last successful attempt, and they can begin piecewise-decrypting the body of the sensitive code itself (e.g., 2432).

[0446] Finally, another virtual join (2430) can be performed, and control can be transferred to the sensitive code (2434) (optionally, after another round of decryption) Again, it is desirable to transfer control to the sensitive code only if the virtual join happens soon enough. If too much time passes, each virtual thread can simply start obliterating the useful information it has placed in memory, and then enter a complex and endless loop of obfuscation statements.

[0447] Additional techniques that are particularly well-suited for use in protecting a sensitive code segment include:

[0448] Inserting genuine race conditions. For example, a few genuine race conditions can be inserted such that multiple attempts by an attacker to follow program execution will get different results. Optionally, termination of loops can depend on relatively infrequent resolutions of the races.

[0449] Inserting disinformation. The obfuscation segments are also a good place to insert disinformation. For example, strings can be manipulated (e.g., print,

concatenate, etc.) with suggestive values (e.g., “debugger detected”, “secret key invalid”, etc.), so as to divert the attacker's attention.

[0450] Decrypting code. Algorithms—including relatively simple algorithms such as ROT13, XOR with “5A5A5A . . .”, etc.—can be used to decrypt code for use in other obfuscation segments.

[0451] Performing system operations. Code can be inserted to, e.g., read the clock or initiate disk operations, although care should generally be taken to avoid things that take too long or that adversely affect the timing of a virtual thread.

[0452] Thus it can be seen that a variety of techniques can be used to elude, detect, or mislead attackers before the execution of sensitive code. Startup and places where the SDP receives or replies to RPCs or makes system calls are likely points of attack and thus make good candidates for protection using these obfuscation techniques.

[0453] 1.6. Data Structure Obfuscation

[0454] 1.6.1. Hidden Invariants

[0455] A program can be rendered effectively opaque by embedding invariants that are difficult for an attacker to infer, yet relatively simple for an obfuscator to maintain. An easy way to embed an invariant is via an opaque predicate. For example, in the code sequence “if P then E1; else E2;” the obfuscator can select P such that it is impractical for an attacker to determine whether P is constant (e.g., always true or always false) or variable. If the attacker does not know that P is constant, he must spend time understanding both E1 and E2.

[0456] One technique for creating opaque predicates makes use of aliasing. For example, determining whether two pointers point to the same location, or into the same structure, is an effectively intractable problem. Opaque predicates involving pointer equality are particularly useful, since the obfuscator will know about the partitioning of the pointer variables (e.g., between ones that point into a first structure and ones that point into a second), but it is not easy for an attacker to infer this information automatically. Substantially the same effect can be achieved with an incrementing counter, or with a pseudo-random number generator with a suitably long period. The important point is that the obfuscator knows, by construction, something that is hard for an attacker to deduce.

[0457] This type of aliasing analysis is effectively equivalent to subscript analysis for arrays. That is, it is typically just as hard to determine whether two subscripted references are to the same element of an array as it is to determine whether two pointers are aliased. Thus, if A is an N-element boolean array, pointer references can be replaced by suitable array references.

[0458] For example, if ve_i and vo_i are variables containing even and odd values, respectively, and N is even, then after the code sequence: “A[ve1%N]=False; A[ve2%N]=True;” the conditional statement “if (A[ve1%N])” could branch either way (depending on the values of ve1 and ve2) since $ve2\%N$ may evaluate to the same value as $ve1\%N$. However, after the code sequence: “A[ve1%N]=True; A[vo1%N]=False;” the conditional statement “if

($A[\text{ve1}\%N]$)” will always take the true branch. (Note that “%” denotes a “mod” operation).

[0459] There are a variety of other techniques for constructing opaque predicates. For example, if the obfuscator has a set of variables that always hold even numbers, and another set that always hold odd numbers, opaque predicates can be constructed from the fact that no even number equals an odd number, that the product of two odd numbers is odd, that the difference between two odd numbers is even, and so forth. Assignments (or procedure calls, etc.) can be sprinkled throughout the program to update these variables in ways that maintain the invariant. For example:

```
[0460]  ve1=ve2+2;
[0461]  vo1=ve1+vo2;
[0462]  ve3=ve1*ve4;
[0463]  vo5=vo6*vo7;
[0464]  ve4=2*vo1;
[0465]  ve2=vo1-vo4+ve1;
[0466]  ve1+=2;
[0467]  ve2*=2; and/or
[0468]  vo4=vo3 % 16;
```

[0469] An advantage of using odd/even variables is that conventional tools are unlikely to analyze variables for oddness/evenness invariants, yet these invariants can be maintained relatively easily with ordinary (e.g., stealthy and cheap) operations. The use of odd/even variables also has the advantage of working with virtually any size of integer, and is preserved modulo any even integer, including, as a special case, integer overflow. It should be appreciated that there are numerous other simple predicates on integers that could be used instead of, or in addition to, the odd/even distinction (e.g., positive/negative, divisible by n , etc.), especially if it became known that deobfuscators should look for variables that are always odd or always even.

[0470] 1.6.2. Data Scrambling

[0471] Another way to obfuscate data structures is to take a number of different data structures—some constant and some variable—and place them in an array, OBA. This technique is illustrated in FIG. 25, in which a variety of data structures and variables 2502 in a computer program are merged into an array 2504 in an obfuscated version of the computer program. For each individual variable (or constant), v , a distinct compile-time constant index, iv , is selected, and the program is appropriately modified to reference $OBA[iv]$ instead of v . Thus, in FIG. 25, for example, a reference to variable i 2510 in the original program can be replaced with a reference to $OBA[1]$ 2511 in the obfuscated program.

[0472] In many cases, OBA can contain arrays as well as scalar values. For example, an n -element array, A , that is indexed by constants can be placed in OBA. Distinct compile-time constant indices, $Ai1, \dots, Ain$, can be selected and used to systematically reference $OBA[Aik]$ instead of $A[k]$. This technique is illustrated in FIG. 25, in which array B 2506 is incorporated into array OBA 2504, and the three elements of array 2506—i.e., $B[1]$, $B[2]$, $B[3]$ —are referenced in array 2504 as $OBA[m]$, $OBA[m+1]$, and $OBA[m+$

2], respectively. The elements of an array in the original program need not appear contiguously in OBA. Thus, in FIG. 25 the elements of array 2508 in the original program are separated in array 2504, with a first portion of the elements 2514 being separated from a second portion of the elements 2516 by one or more other variables, constants, or elements from other arrays (represented in FIG. 25 by ellipses 2515). Similarly, even when elements of an array do appear together in OBA, they can be rearranged so as to confuse or obscure the logical relationship, if any, between the elements in the original array. Thus, as illustrated in FIG. 25, the order of elements 2516 can be chosen so that it does not correspond to the order in which the same elements appear in the rows and columns of array 2508. If an array in the original program is indexed only by simple loop variables, the loops can be unrolled to get rid of the variable indices in the obfuscated program. If an array is indexed in more complex ways, it can be allocated to sequential locations in OBA, or it can simply be excluded from this scheme.

[0473] At SDP customization time, the elements of OBA can be permuted, respecting array elements that must be contiguous. To accomplish such a permutation, constants allocated within OBA are moved, and code that indexes into OBA is patched by replacing the compile-time constant index with the result of the permutation. The execution of the program will be unchanged, but differently-customized instances of the SDP will have different storage layouts, thus making them more difficult to recognize. In addition, since arrays will be dispersed, they will be harder to identify.

[0474] Other data structure permutations are also possible. For instance, several arrays might be merged together, and their indexes adjusted accordingly. For example, the even index elements might belong to logical array, A , and the odd index elements to logical array, B . Moreover, the mapping between arrays and index locations can be varied for different program customizations.

[0475] 1.7. Optimization as Obfuscation

[0476] Program optimization can be viewed as a process of systematically sacrificing program structure to gain performance. As a result, optimized code is typically more difficult to debug or understand than code that is compiled in a more straightforward fashion. Thus, many optimizing transformations also tend to be good obfuscation transformations.

[0477] Optimization may involve moving computations around, combining variables, changing the flow of control, replacing straightforward but expensive computations with complex but cheaper ones, and introducing and exploiting hidden invariants, to name just a few of the processes that are commonly performed during optimization. In general, optimizing transformations:

[0478] Have varying potencies, depending largely on how much context they exploit. For example, peephole optimizations typically have low potency; local optimizations tend to have only moderate potency; global optimizations generally have higher potency; and interprocedural optimizations can be extremely potent. They are all, of course, stealthy.

[0479] Are highly resilient. Since optimizations typically remove structure rather than add it, it is difficult

to reverse these optimizations automatically. Indeed, many optimizations are one-way.

[0480] Have low cost. From a performance perspective, optimizations tend to have a negative cost; that is, they typically improve performance rather than degrade it.

[0481] One difficulty in performing program optimization is deciding whether a given transformation will improve program performance or degrade it. However, for obfuscation purposes a transformation can typically be applied without giving much consideration to whether it increases program performance or decreases it, as long as the transformation does not decrease performance too drastically.

[0482] Accordingly, techniques for program optimization, including without limitation constant folding and propagation; peephole transformations; elimination of common sub-expressions and other redundant code; loop unrolling and procedure in-lining (especially if different program instances are customized differently); strength reduction; hoisting and other code motion techniques; local reordering of instructions; partial evaluation; allocation of registers to variables; and other suitable techniques make good candidates for obfuscation.

[0483] Since a good optimizer may undo some of the work done by obfuscation transformations that are designed to add complexity to a program, it will often be more effective to apply these obfuscation transformations after the optimizer has been applied to the program.

[0484] 1.8. Performance and Properties

[0485] While it is generally undesirable for obfuscation to seriously degrade the performance of a program by, e.g., performing a substantial amount of additional computation, it should be understood that there are places where it can be advantageous to perform extra computation. For example:

[0486] While trying to lose an attacker. As discussed above, it is desirable to impede an attacker's efforts to track the execution of sensitive regions of code. By employing the techniques described above, the amount of time an attacker takes to finish executing sensitive code can be increased. Since an attacker—unlike a normal user—is likely to want to repeatedly execute the sensitive code, the waste motion and delay added by the extra obfuscatory prologue computations are likely to present a non-trivial obstacle to an attacker's attempts to run experiments on the program.

[0487] While dealing with external devices. Another time in which heavy-duty obfuscation can be used at relatively little cost is while dealing with external devices, since the cost of the obfuscation will often be overshadowed by device and controller latency.

[0488] Thus, if an application is performance-sensitive, obfuscation techniques can be concentrated in certain parts of the program to lessen the effect (if any) that obfuscation has on the program's performance.

2. Tamper Resistance

[0489] An attacker should find it relatively difficult to tamper with the operation of an SDP, and should find it especially difficult to do so undetected. To this end, a tamper resistant program should possess some or all of the following characteristics:

[0490] Attacks on an intact SDP (e.g., running the SDP under a debugger, patching code, etc.) should be detected promptly—or at least made more difficult to perform undetected.

[0491] Attempts to run parts of an SDP—as opposed to a complete, properly-installed SDP—should generally be detected and/or rendered ineffective.

[0492] Suspected tampering should be reported externally for fraud detection, response, and/or prevention as soon as practical. In addition, in a preferred embodiment pressure to connect to an external site is subtly increased once tampering has been detected. For example, in an application in which a user is allotted a certain budget with which to perform various offline activities, the budget can be made to expire when tampering is detected (or soon thereafter), thus forcing the user/attacker to connect to an external site that preferably houses, or is communicatively coupled to, the external tamper-monitoring agent. It will be appreciated that other techniques for encouraging the user to connect to an external agent could be used in accordance with the principles of the present invention and the requirements of the particular application.

[0493] After tampering has been detected, critical information residing in the SDP (e.g., master keys, critical code segments, etc.) should be erased or modified relatively quickly. Preferably, however, this should not occur until after tampering has been reported and/or other appropriate defensive action has been initiated, since the SDP will generally need to be functioning in order to send a report to an external agent, and since it is desirable to avoid alerting the attacker to the fact that an attack has been detected for as long as possible, and preferably at least until other defensive action has been initiated or completed. To protect against false alarms, in a preferred embodiment only information that can be restored by, e.g., the external agent(s) is erased, while other information is left generally intact.

[0494] Useful functioning of the SDP should be restricted or prevented once tampering has been detected. In one preferred embodiment the SDP operates with limited functionality, rather than stopping entirely, so that the value to the attacker of continuing to run the program is lessened, and so that the attacker is not alerted to the precise location at which detection occurred.

[0495] Tamper resistance mechanisms should be relatively independent of the design of the SDP, thus minimizing the need for those involved in implementing the SDP to communicate with those involved in implementing the tamper resistance controls.

[0496] To accomplish these goals, preferred embodiments of the present invention make use of some or all of the following:

[0497] Tamper Testing. Tamper testing generally involves looking for evidence that the SDP may be under attack. For example, the tamper testing mechanism can check to see if the SDP is connected to a debugger, has code that has been modified (e.g., breakpoints set or instructions altered), and/or has key data structures that have been improperly changed.

[0498] **Tamper Memory.** Tamper memory is used to robustly record the fact that one or more tamper tests have detected suspicious activity, at least until certain tamper reaction measures have been initiated or completed.

[0499] **Tamper Reaction.** Tamper reaction generally refers to the process of reporting the detection of suspicious activity to an external agent and/or taking steps to limit the damage that a tampered SDP, or an attacker, can do.

[0500] 2.1. Tamper Testing

[0501] In general terms, tamper testing involves looking for evidence that the SDP is under attack. For example, tamper testing can involve looking for evidence that an SDP is connected to a debugger, has been improperly modified (e.g., breakpoints set or instructions altered), and/or that key data structures have been improperly changed. To this end, the debugger-detection, watchdog, and self-validation techniques described herein can be used. In addition, an SDP can use, e.g., operating system calls to keep track of changes to the SDPs operating environment, such as changes to the modules loaded in system memory (e.g., the removal or addition of modules), changes to interrupt or exception handlers, and the like.

[0502] In a preferred embodiment tamper resistance and obfuscation measures are applied in a relatively pervasive (as opposed to local) fashion, thereby reducing the chances that an attacker will recognize the protection mechanisms that are being used and/or find and eliminate each test for tampering. While pervasive application of obfuscation and tamper resistance techniques can be impractical in certain situations, the use of such techniques can be directed or limited to certain areas of the SDP to the extent that there is a material engineering or support cost to applying them. For example, if source code must be manually obfuscated or instrumented for tamper-resistance, that would argue in favor of applying tamper resistance in a way that minimizes such costs. Thus, in one embodiment performance-critical regions are flagged, and tamper testing code is not added to the flagged regions. In a preferred embodiment conventional automated tools are used to identify and flag the performance-critical sections, while in other embodiments, other techniques are used (such as, e.g., manual identification).

[0503] 2.1.1. The Priming-Stamping Approach

[0504] Many tamper-resistance mechanisms can be implemented using a relatively small collection of macros and a binary image postprocessor. The macros can be transparently inserted into the code without substantially changing the semantics of the code and without requiring the user to know much, if anything, about the macros' implementation. The binary image postprocessor finds where the macros were used by looking for the distinctive primers that were created by the macros (preferably post SDP-customization), and implements tamper-resistance mechanisms in the spaces reserved by the primers. This technique is illustrated in FIG. 26.

[0505] In one preferred embodiment, some or all of the following macros are used:

[0506] **IMPORTANT_CODE:** As shown in FIG. 26, this macro **2610** is preferably inserted relatively close

(e.g., within a thousand bytes) to code **2602** that is critical to the correct functioning of the SDP **2600**. It serves as an indication by the programmer that a self-test **2620** should be performed at runtime to ensure that the nearby code **2602** has not been altered. Although a programmer will generally need some knowledge of the system structure to insert macro **2610** optimally, insertion can be done quite easily and precise placement is typically unnecessary. As shown in FIG. 26, a binary image postprocessor can be used to find the primer **2611** created by macro **2610** in the binary image **2600b** of program **2600a**. The desired self-test code **2620** is then inserted into the binary image in the space reserved by the primer, and executed when program **2600** is run.

[0507] **TEST_FOR_TAMPERING:** This macro **2612** can be inserted in a number of places in the source code **2600a** to trigger full or partial tests for tampering **2622**. For purposes of practicing the present invention, any suitable tamper testing code **2622** can be inserted. For example, tamper testing code **2622** may include code that checks for the presence of a debugger, checks for modification to the code **2600b**, and/or for other suspicious activity. Macro **2612** is preferably used in places that are likely to be executed repeatedly—so that tests are done often enough—but preferably not in inner loops, where the cost of tests might become an issue. Again, some knowledge of the system structure will typically be needed to insert this macro optimally, but once again, it can be done quite easily and precise placement is generally not important.

[0508] **TAMPERING_DETECTED:** This macro **2614** is operable to trigger the insertion of code **2624** that returns a boolean value that is true if the tamper memory shows evidence of tampering. Macro **2614** is preferably called in places **2604** where the SDP is about to access sensitive or valuable information, or where there is an opportunity (or need) to “nudge” the user to connect to an external site to enable tamper reporting. Since deciding where to test **TAMPERING_DETECTED**—and how to respond when it returns true—generally requires knowledge of critical parts of the SDP **2600**, it is preferably inserted by the programmers of the SDP source-code **2600a**.

[0509] **REINFORCE_TAMPER_MEMORY:** This macro **2616** is preferably inserted in a number of places in source code **2600a** to provide positive feedback in the tamper memory mechanism **2626**. It should appear in places that are likely to be executed repeatedly, so that amplification is done often enough. Since it is typically not as expensive as **TEST_FOR_TAMPERING** **2612**, it can be placed in more frequently-executed code. As described in more detail below, this macro can be used to implement the “watchdog” scheme **2626**. In general, its roles are (i) to obscure the connection between instances of **TEST_FOR_TAMPERING** and **TAMPERING_DETECTED**, thus making it harder to find and disable the tests, and (ii) to create redundant paths between instances of **TEST_FOR_TAMPERING** and **TAMPERING_DETECTED**, thus making it more likely that even if a path is found and disabled, other paths will still function. Again, some knowledge of the system structure will typically be needed to insert this

macro optimally, but it can be done quite easily and precise placement is generally not important.

[0510] In one embodiment, TEST_FOR_TAMPERING 2612 and/or TAMPERING_DETECTED 2614 include IMPORTANT_CODE 2610, since self-testing is typically an important target for tampering. In such an embodiment TEST_FOR_TAMPERING will thus typically check the validity of more than one block of code.

[0511] 2.1.2. Self-Validation

[0512] One scheme for self-validation involves computing a checksum (or a secure hash) for a region of code before the code is executed, and comparing the computed checksum to a previously stored checksum value for that code region. If the checksums match, the code is deemed to be valid; otherwise, appropriate defensive action is taken. As another example, space can be reserved in the binary image of a program for a “correction” bit-pattern that makes the checksum for a particular region come out to a predefined value, e.g., zero. When a code block is to be executed, the checksum for the region can be computed and compared against the predefined value. If the calculated checksum does not match the predefined value, this can be used as evidence that the code in the original program has been modified, and appropriate defensive action can be taken. In a preferred embodiment, the CS prefix is used to ensure that checking is performed on the code that is actually being executed.

[0513] In one embodiment, none of the checked regions contain checksums for other checked regions. However, this generally means that the checksums themselves will not be checked for tampering. Thus, as illustrated in FIG. 27, in another embodiment the checksum 2714 for a region of code 2704 is embedded in another region of code 2706, and a checksum 2716 for code region 2706—including embedded checksum 2714—is computed and stored in a third region 2708. This process can be repeated for more regions of code (e.g., 2702, 2710). The final checksum value(s) 2720 are preferably encrypted, cryptographically signed, stored in tamper-resistant memory, and/or otherwise protected from undetected alteration.

[0514] Care should be taken when inserting the checksums or correction values for a given code region to avoid altering the checksums of other code blocks, since modifying a code block without appropriately modifying its corresponding, stored checksum will cause self-validation of that block to fail. Accordingly, in one embodiment checksums for overlapping code regions are set using a directed graph that represents the checksum dependencies, as shown in FIG. 28. Each vertex 2810 represents a checked region. Each edge 2820 goes from a region 2810 containing a checksum 2812 to the region to which the checksum corresponds. If this graph is cyclic, as shown in FIG. 28B, then it will generally not be possible to make all the checksums mutually consistent. However, if the graph is acyclic, as shown in FIG. 28A, then it will be possible to make the changes (e.g., insert the checksums and/or correction bits) in an order consistent with the graph, for example, by using a topological sort.

[0515] Another relatively simple technique is to store checksums in the data section, thus avoiding an effect on the checksums for the code section. A checksum can be computed separately for the portion(s) of the data section containing these checksums and compared against a correct value stored outside the checksummed range.

[0516] Thus, it will be appreciated that self-verification can be designed so that it is relatively complicated to untangle—with checks in one part of the program validating checks in other parts of the program—so that the checking code cannot be patched out easily. Self-modification techniques, such as those described elsewhere herein, may also be used to provide additional protection.

[0517] 2.1.3. Telomeres

[0518] Another tamper testing technique can be thought of as a software analog to telomeres, which are believed to be the genetic mechanism that keeps cells from reproducing indefinitely. An illustration of this technique is shown in FIGS. 29A, 29B, and 29C.

[0519] As shown in FIG. 29A, a few extra blocks of “wallpaper” 2904 are included at the end of SDP executable file 2900. For example, wallpaper 2904 may consist of non-functional code and/or unused data, preferably disguised to look similar to code or data that is actually used by the program. As shown in FIG. 29B, when the SDP startup code 2902 is run, it overwrites the last block 2906 in the executable file 2900b with a predefined pattern, such as a string of zeros, alternating ones and zeros, or the like. Preferably much later, SDP 2900b checks to see if the write succeeded (by, e.g., checking for the presence of the pattern) and takes defensive action if it did not. Still later, preferably after the startup sequence has finished executing and the SDP has performed a variety of tamper tests, the SDP restores wallpaper 2904 to the overwritten block(s) 2906. If startup code 2902 is executed again before the wallpaper is restored—as might happen if an attacker terminates program execution and restarts the program—the wallpaper in block 2905 will also be overwritten with the predefined pattern. For example, in one embodiment startup code 2902 contains code 2903 operable to overwrite the first block at the end of the program (or other suitably predefined location) that has not already been overwritten by the pattern. As shown in FIG. 29C, if there are too many consecutive attempts to restart the SDP, some critical executable code will get overwritten, and program execution will fail. For example, as shown in FIG. 29C, if startup code 2902 is restarted three consecutive times, pattern 2909 will overwrite genuine code block 2908, and execution of program 2900c will thus fail when pattern 2909 is encountered. The above-described telomere techniques are preferably used as a supplementary security measure in combination with other security measures.

[0520] 2.2. Tamper Memory

[0521] In a preferred embodiment, tamper memory is provided to record information regarding suspicious activity so that the SDP and/or an external agent can later respond appropriately. Tamper memory should be robust; that is, it should be relatively difficult for an attacker to locate and/or modify or circumvent the tamper memory mechanism(s).

[0522] 2.2.1. Watchdogs

[0523] One tamper memory mechanism makes use of “watchdogs” that are distributed (and possibly hidden) throughout the program, the watchdogs being operable to record evidence of tampering. For example, in one embodiment watchdogs are scattered throughout the system’s global memory (e.g., from 0.1% to 10% of the total). In a preferred embodiment the watchdogs are simply extra vari-

ables in the SDP's memory, each of which may have values associated with the interpretations "quiet" and "barking" (i.e., "tampering detected" and "tampering not detected," respectively). Each variable may be of virtually any type, and its two values may be encoded using virtually any predefined distinction. For example, if a boolean variable is used, false could correspond to quiet, and true could correspond to barking, or vice versa. Similarly, for an integer variable the distinction could be between zero and non-zero, positive and negative, odd and even, and so forth. A character variable could distinguish between alphabetic and numeric values, between vowels and consonants, between printing and nonprinting, etc. As yet another example, a pair of variables could represent quiet by being different and barking by being the same, or vice versa. Thus, it will be appreciated that any suitable scheme for differentiating between quiet and barking (and/or additional or equivalent states) may be used.

[0524] FIGS. 30A, 30B, and 30C illustrate the use of watchdogs. As shown in FIG. 30A, when a program fails a tamper test (e.g., evidence of a debugger or similar tool is detected, or alteration of a piece of code or watermark is detected) (3002), the program puts barking values in a predefined number of watchdogs (e.g., a few) (3004), but may take no further defensive action at that time.

[0525] However, as shown in FIG. 30B, before the program does something that might be valuable to the attacker (3014), the program tests a predefined number of watchdogs (e.g., a few) (3010), not necessarily in a contiguous code sequence. If one or more of the tested watchdogs is barking, the program can initiate a suitable tamper response (3012). For example, the program may do something that is less useful to the attacker, and/or that enables external reporting, such as indicating that the attacker needs to connect to an external site for an upgrade. In another embodiment, when barking watchdogs are detected the program simply transfers control to interminable obfuscation code that looks like real code, but actually makes no useful progress.

[0526] In one embodiment, short sequences that test a random watchdog are scattered throughout the program (i.e., are not limited to potentially sensitive code segments). As shown in FIG. 30C, if one of these sequences determines that a watchdog is barking (3020), it can place barking values in a predefined number of other watchdogs (e.g., a few) (3022), thus spreading barking throughout the code and providing a form of positive feedback.

[0527] Note that, as used above, "a few" watchdogs refers generally to any quasi-random subset of preferably at least two (and in one embodiment not more than four) watchdogs chosen differently at insertion time for selected points in the code. This sort of randomization can be used to complicate the task of finding all of the watchdogs and all of the tests, while increasing the likelihood that barking will spread quickly and widely once initiated.

[0528] When the system has an opportunity to communicate with an external site, the system can send a block of memory that includes several watchdogs (without necessarily testing the watchdogs itself) for analysis by an external fraud-detection process. For example, in round-trip exchanges an external site can request that a range of the SDP's memory be returned. The external site can inspect, inter alia, the watchdogs contained in the requested memory

range (using, e.g., a key that indicates where the watchdogs are located), thus giving away little or no information to an attacker about where the watchdogs are actually stored.

[0529] It should be appreciated that watchdogs need not be represented by single variables, or single bits, in memory. In general, they can consist of any predicates representing invariants on (correct) system states. In addition, in one embodiment the mapping from quiet/barking to true/false (or other appropriate values for non-boolean variables) is determined by an insertion-time random mask. Individual watchdogs can be distributed across multiple variables to ensure that the values stored in memory do not remain constant, and that replacing load instructions by immediate values will not keep a watchdog quiet (unless all loads from its components are found and patched at the same time).

[0530] In addition, barking propagation sequences can be inserted in place of, or in addition to, null-effect sequences during obfuscation, instance customization, or watermarking. Barking propagation sequences have the advantage that since they are not literal NOPs, and do affect the program state, they are generally more difficult for an attacker to recognize than true null-effect sequences. Shuffling the components of a distributed watchdog is another good NOP candidate. For example, if a watchdog is represented as the XOR of two boolean variables, it is effectively a NOP to, e.g., swap the values of the two variables, complement the values of both variables, or assign the XOR of the two variables to one of the variables and assign zero to the other. Since these operations may change the values of variables in memory, they generally cannot be recognized as NOPs unless the implicit relationship between the variables is known.

[0531] The selection of watchdogs can also be used to encode information about the identity of the instance. For example, the selection of watchdogs can serve as a form of watermark. Thus, it will be appreciated that the use of watchdogs, like the use of other techniques described herein, need not be limited solely to playing a role in performing one function (e.g., tamper resistance, obfuscation, or watermarking), but can instead play a role in accomplishing a combination of these functions (e.g., obfuscation, tamper resistance, and watermarking).

[0532] It should be appreciated that watchdogs can be set based on the occurrence of other conditions besides the failure of an explicit tamper test. For example, a calculation can be performed, and its results stored in such a way that if something is wrong, a watchdog starts barking. The redundant tests described above are an example of such implicit tests. Similarly, testing a watchdog need not involve an explicit test in the code, but can instead consist of, for example, performing a calculation that yields a different outcome if the watchdog is barking.

[0533] Placement of the watchdogs can vary from instance to instance of the SDP. For example, a larger pool of bits can be reserved, from which a random subset is drawn at insertion time.

[0534] Use of a watchdog mechanism (or any similar mechanism), generally makes it more difficult for an attacker to identify and disable tamper testing in the SDP. In addition, because it is independent of particular methods of tamper testing, it can be used in combination with virtually any

collection of tamper testing techniques. Moreover, by effectively separating the tamper testing and the tamper reaction, and by inserting positive feedback into tamper memory, this mechanism makes it more difficult for an attacker to evade detection.

[0535] 2.3. Tamper Reaction

[0536] Once tampering has been detected, any of a variety of techniques can be used to respond. For example, it is generally preferable to make the attacker think that he is making progress for as long as practical under the circumstances, and to make it relatively difficult for him to determine where he went wrong. Thus, as shown in FIG. 31A, in a preferred embodiment if tampering is detected (i.e., a “yes” exit from block 3102), the SDP branches to disinformation code (3104), and also preferably spoils locally-held keys and/or records information about the attack for eventual transmission to, or detection by, an external fraud detection agent (EA). Another stratagem is to start simulating disk or other I/O errors (3106), so that the attacker will think his problems are due to poor hardware and/or a poor operating system. Another tamper reaction is for the SDP to simply shut down immediately, although this may give the attacker clues regarding the location of the tamper detection mechanism.

[0537] Yet another tamper reaction mechanism targets the people who actually crack the SDP. Features are built into the SDP that selectively cripple cracked systems without affecting legitimate systems. For example, when there is evidence that the SDP is under attack or has been successfully attacked, key operations can be made to fail mysteriously and seemingly non-deterministically, as though the operating system were reporting a variety of errors (e.g., disk read errors, out of virtual memory errors, memory protection violations, execution of illegal instructions, etc.) at a low enough rate not to implicate the SDP, but at a high enough rate to make it difficult for the attacker to continue using/attacking the SDP. This form of tamper reaction is especially appropriate for systems that connect to external sites infrequently if at all, although it will be appreciated that it is readily applicable to other systems as well.

[0538] FIG. 31B illustrates another exemplary tamper reaction mechanism. As shown in FIG. 31B, upon receiving a return code after an operating system call (3110), a watchdog is tested (3112). If the watchdog is barking (i.e., a “yes” exit from block 3112), a random number can be generated (3114) and compared to a predefined threshold (3116) to decide whether to branch to the code that would be executed if a real error code had been returned by the operating system call (3118), or whether to proceed with normal execution (3120).

[0539] 2.3.1. Forcing the SDP to Validate Itself

[0540] It will often be desirable for external agents to force an SDP to prove its own validity. FIG. 32 illustrates one such technique. Referring to FIG. 32, a table 3202 of random keys 3203 (e.g., 1000 or some other suitably large number) is created when an SDP binary 3210 is built. Keys 3203 are used to calculate hash values 3204 for various regions 3206 of SDP 3210. For example, a MAC-type hash function can be used to compute fixed-sized “digests” 3204 of arbitrary-sized blocks of code and/or data 3206, such that it is relatively difficult to find another block 3206 that yields

the same MAC value. It will be appreciated that there are a variety of techniques for calculating such a MAC or hash value, any suitable one(s) of which can be used, including those set forth in Menezes et al., “Handbook of Applied Cryptography,” pp. 321-83 (CRC Press, 1997), which is hereby incorporated by reference.

[0541] External agents 3200 can keep a copy of table 3202, and when an agent 3200 wishes to validate a particular instance of the SDP 3210, agent 3200 can send a key 3220 to the system 3212 running SDP instance 3210, and tell it to calculate the MAC value for the appropriate code region 3222. System 3212 uses key 3220 to calculate a MAC value 3224, which it sends to agent 3200. Agent 3200 compares received MAC value 3224 with the appropriate MAC value 3204 in table 3202. If the MAC values match, then the SDP is deemed valid; otherwise, appropriate defensive action is initiated (such as one or more of the tamper reaction techniques described herein). It will generally be difficult for an attacker to know the correct MAC value to return if the requested code region 3222 has been tampered with.

[0542] For robustness, the hash values 3204 are preferably relatively large. In addition, in a preferred embodiment a relatively large number of possible validation queries are defined (e.g., the size of table 3202 is relatively large). Additional security can be provided by sending system 3212 a program 3226 to calculate MAC 3224, in addition to sending key 3220, thus making it much more difficult for an attacker to fool external agent 3200 by tampering with the MAC calculator than would be the case if the MAC calculator formed part of SDP 3210 or was otherwise permanently stored on system 3212.

[0543] For additional security, code self-modification techniques (such as those described herein) are preferably used in SDP 3210, such that SDP 3210 modifies itself from time to time in ways predictable by external agent 3200 (e.g., copying the current date into a code location). As a result, queries 3219 directed to an out-of-date version of SDP 3210 will generate incorrect answers which external agent 3200 can detect, even if the original version of SDP instance 3210 continues to run. In addition, in a preferred embodiment the SDP can have a mechanism that only allows one instance of itself to be run at any one time on a given machine 3212.

[0544] The MAC calculation code contained in the SDP and/or included in query 3219 can also be written to have a stealthy side-effect (e.g., setting a watchdog-like variable) in addition to reporting results to the external agent 3200. The presence of the side-effect can be tested in subsequent SDP code and/or reported to an external agent in a subsequent interaction. For example, in one embodiment the stealthy side effect includes setting a flag. The flag is cleared if it is not reset within a predefined period of time. The SDP (or external agent) checks the flag periodically, and if it detects that the flag has been cleared, it can deduce that the SDP has not been asked to validate itself for at least the predefined time period. This may indicate that someone has disabled the self-validation mechanism or otherwise tampered with the system. Additional self-validation techniques are described in the Shear ’754 application and the Sibert application, each of which was previously incorporated by reference herein.

[0545] 2.4. Time Management

[0546] Many programs are largely dependent on the host platform’s time services, which can often be manipulated

with relative ease to change the platform's notion of date and time and thus potentially subvert the intent of time-related mechanisms used by the program. This section describes some useful techniques for maintaining an accurate time measurement, including:

- [0547] Adjustment for the difference between absolute time and the instance's time (which may become unsynchronized because of clock drift or operator error, entirely without malicious intent);
- [0548] Determination of the correct local time (e.g., time in the correct time zone) for the SDP instance, and optionally for each user of the instance;
- [0549] Automatic update of the time and time zone information during communications with external agents; and
- [0550] Countermeasures to detect and defeat deliberate attempts to adjust the SDP instance's time values.

[0551] Additional time management mechanisms are described in the Ginter '900 patent, previously incorporated by reference herein. In addition, in a preferred embodiment obfuscation and tamper resistance techniques, such as those described above, are used to, e.g., obscure and/or protect the operation of the time management mechanisms.

[0552] 2.4.1. Time Offset Mechanism

[0553] A basic time management mechanism uses a signed time offset that relates the platform's time to external time (e.g., indicates how much the SDP platform's system clock differs from a predefined external time source). As shown in FIG. 33, in a preferred embodiment this offset 3304 is stored with each SDP instance 3302, and is retrieved at instance initialization and added to the platform's time value 3306 whenever a time reading is required by SDP instance 3302. Offset 3304 can be determined initially and/or each time a connection 3308 is made to a trusted external source or agent 3310 by checking the platform's clock 3306 and calculating the difference between that time and the external time 3312 obtained from external source 3310. It will be appreciated that if there are multiple external sources, their times can be synchronized using conventional synchronization techniques. For example, techniques such as those described in Lampart et al., *Synchronizing Time Services*, DEC/SRC Report 18 (June 1987) can be used.

[0554] Once the correct offset 3304 is determined, this technique ensures that SDP instance 3302 knows what time it is, as long as the SDP platform's clock 3306 does not drift and is not explicitly changed. Of course, these ideal conditions will typically not apply, since under normal conditions clocks can be expected to drift, to be explicitly set or changed, and so forth. Thus, an update mechanism such as that described below is preferably used to keep the SDP's time offset 3304 relatively synchronized with external time 3312.

[0555] 2.4.2. Time Zone Correction

[0556] In addition to the time offset, in one embodiment a time zone offset is maintained by the SDP instance, the time zone offset being operable to facilitate conversion of external time to local time. During instance initialization, the user can be queried for the user's time zone, which can also be used as the time zone of the SDP instance. Based on this

information, the SDP can compute the zone offset and store it in its database along with the time offset.

[0557] Unlike the time offset, however, which is per-instance, the zone offset may be different for each user of the SDP instance. In the case of a multi-user instance, for example, different users may be in different time zones. Accordingly, at user activation the user's time zone can be requested and used to calculate the appropriate time zone offset for that user.

[0558] Moreover, because of daylight savings time (and similar mechanisms throughout the world), the time zone offset for a particular user may change as time passes. Although these changes are relatively infrequent—e.g., usually twice per year—it is desirable to accommodate them automatically. Unfortunately, there is no simple rule by which such changes can be made, since they are subject to the vagaries of local governments. Thus, as shown in FIG. 33, in one embodiment the time zone correction is expressed as a list of pairs 3320, each pair including the following information:

[0559] Zone Offset. Zone offset 3322 specifies the offset for calculating local time.

[0560] Zone Change Time. Zone change time 3324 specifies the external time at which the corresponding zone offset changes to the zone offset in the next pair in the list.

[0561] The SDP 3302, or one or more of the external agents 3310, can calculate these values at instance initialization, and supply a list 3326 that is valid for a period of several years forward and backward. Whenever the SDP instance is started, it can look at list 3326, find the appropriate offset 3322 for the current value of external time, and set a background manager job to run at the next zone change time 3324 and to adjust the zone offset appropriately. Of course, on most days this background job will not run, nevertheless, it is preferably set each time to ensure that the zone offset update occurs when needed during normal SDP operation.

[0562] If the user switches time zones, a different set of time zone information will be needed. This can be provided at activation time (if the user supplies a list of possible time zones), or can be updated later. For example, in one embodiment an SDP instance is operable to store sets of time zone information for different time zones in order to allow a user to switch time zones without interacting with an external agent. Since the SDP instance knows the absolute time via the time offset, a user will generally be unable to tamper with the SDP by changing the time zone, since the time zone data simply affects the time display, not the SDP's internal functionality.

[0563] It will be appreciated that there are a variety of ways to store the time-related information discussed above, any of which can be used without departing from the principles of the present invention. FIG. 33 illustrates one exemplary embodiment. Referring to FIG. 33, in one embodiment SDP 3302 is operable to store a Time_Offset record 3304 in a per-instance configuration variable table 3330, and to store a Zone_Info record 3332 in each user's configuration variable table 3334 as well as the per-instance configuration variable table 3330. Time_Offset record 3304 holds the offset between external time 3312 and the plat-

form's time **3306**, while the Zone_Info record **3332** holds one or more zone information lists **3326**, and an indication **3338** of the last time the Zone_Info record **3332** was updated by an external agent **3310**. As discussed above, each zone information list **3326** can contain a zone name **3340** and one or more pairs **3320** of time and zone correction values (**3324**, **3322**, respectively).

[0564] 2.4.3. Time Update Mechanism

[0565] Because the platform's clock may drift or otherwise be adjusted, the SDP's time offset should be updated occasionally to maintain synchronization with the trusted external time source. In one embodiment this update occurs automatically whenever an SDP instance initiates an on-line communication with a trusted external agent using a pre-defined communication protocol (CP). This technique is illustrated in FIG. 34.

[0566] Referring to FIG. 34, each CP message **3410** from an external agent **3402** is time-stamped (**3412**) with the external agent's time. Upon receiving the time-stamped message (**3414**), SDP instance **3404** compares the external agent's time with its own time (**3416**). If the difference between the external agent's time and the SDP instance's time is greater than a predefined amount, N (e.g., thirty seconds) (**3417**), the SDP instance automatically updates its current time offset (**3418**), stores it in its database, and uses the new value to calculate the correct time value from the platform time.

[0567] Similarly, when SDP **3404** sends a message **3420** to external agent **3402**, the SDP is operable to stamp the message with the SDP instance's time (e.g., the platform time combined with the SDP instance's time offset, and/or the platform's time and the time offset individually) (**3422**). Upon receiving the time-stamped message (**3424**), agent **3402** compares the SDP instance's time with its own time (**3426**). If the difference between the SDP instance's time and the agent's time is greater than a predefined amount, M (e.g., thirty seconds) (**3427**), and/or if the value of the instance's time offset is different from the last offset value that the agent received from the SDP instance, the agent initiates appropriate defensive action. For example, the agent may record the detection of a change in the instance's time and/or time offset (**3428**), and if a predefined number of such changes are detected in a predefined time period, the agent can, e.g., refuse to grant the SDP permission to perform certain actions, send a message to the SDP that is operable to terminate operation of the SDP or to trigger defensive actions (such as the tamper reaction mechanisms described above), and/or perform any other suitable defensive action (**3430**). Although a time and/or offset update can potentially occur with every CP message **3420**, once an adjustment occurs it is unlikely that normal operation will cause the time to become unsynchronized by more than, e.g., thirty seconds during a single CP session, and thus multiple updates per session can be detected and handled appropriately.

[0568] 2.4.3.1. Offset Adjustment and Monotonicity

[0569] Simply updating the time offset (e.g., replacing an incorrect offset value with the correct value) may cause the SDP to perceive external time as moving backwards (i.e., the SDP may calculate a value of external time that is less than the value it previously calculated). Although updates that

compensate for normal clock drift will generally not change the SDP's perception of external time significantly, updates that correct for deliberate changes or for fraud may result in wildly different computed external times for two apparently adjacent SDP operations. Although, in principle, this can be avoided by using sufficient tamper-resistance mechanisms, the system-level consequences may be relatively unimportant even if it cannot be prevented entirely.

[0570] However, it is generally desirable for routine adjustments (e.g., the thirty second adjustments discussed above) to take place without running time backwards. In one embodiment this is accomplished using a mechanism similar to the BSD UNIX adjtime() system call. The desired time offset is supplied and the real time offset is adjusted gradually until it is synchronized with the correct time, while ensuring that during this process the SDP instance's time service does not return a time value earlier than the previously returned value.

[0571] 2.4.3.2. Time Zone Information Update

[0572] Because the initial time zone information will generally become outdated (e.g., it may be rendered obsolete by governmental action), it is important to have a mechanism by which it can be updated. For example, in one embodiment the SDP instance can refresh its time zone information by communicating at periodic intervals with an external agent. It would be appropriate to make such an update, for example, every time the zone info indicates an impending zone offset change (that is, about twice each year), or, depending on the situation, much more frequently.

[0573] 2.4.4. Sources of Time Change

[0574] An SDP's attempts to maintain accurate time may face some or all of the following challenges:

[0575] Clock drift. The platform clock may not run at precisely the correct rate. This can be compensated for by, e.g., the update mechanism described previously.

[0576] Adjustment for daylight savings time. The platform clock may be adjusted automatically by the host operating system to change when daylight savings time changes, as typically occurs on platforms that do not maintain a zone offset (e.g., Windows). This, too, is easily compensated for by the update mechanism described above.

[0577] Explicit use of Date and Time commands or other time-adjustment interfaces. The user may make explicit use of operating system commands or the user interface to change the time. The user may do this either for normal adjustment (e.g., the user notices clock drift, and corrects it) or in an attempt to commit fraud. The former case can be handled by the update mechanism, described above, while the latter case is preferably handled in the manner described in more detail below.

[0578] These sources of time change apply to the time offset mechanism, as that is where the SDP instance has an explicit dependency on the platform's time services. In general, threats specific to the zone information are not as significant, as the zone information is maintained internally to the SDP instance and is not dependent on platform mechanisms.

[0579] 2.4.5. Countermeasures

[0580] This section describes several mechanisms for handling the sources of time change set forth above. These mechanisms include techniques for updating the instance's time offset, detecting backward time motion, intercepting system interfaces for time-setting, using private instance-maintained time, internally detecting time anomalies, forcing notification of an external agent in suspicious circumstances, analyzing time change patterns, and performing immediate verification of time.

[0581] Most of these mechanisms are operable to detect a single platform time change event. That is, they effectively notice that the time is now $X \pm \delta$ when it was expected to be X . Such changes may be benign (e.g., changing the system clock by a few minutes to compensate for clock drift) or may indicate attempted fraud (e.g., setting the system time back to yesterday to avoid the expiration of a time-limited right). Systems should apply heuristics to decide whether any particular time change is suspicious (e.g., comparing the magnitude of the time change with a pre-defined threshold). Appropriate heuristics will generally depend on the system/application. The heuristics may be applied independently to each event, or may be part of a larger pattern analysis activity performed at the instance, at the external agent(s), or at both.

[0582] Time Offset Update. As described above, a simple mechanism can be used to ensure that an SDP instance's notion of external time is synchronized with real external time each time an interaction with an external agent occurs. While not offering absolute protection, this mechanism is generally effective in ensuring that simple attempts to commit time fraud are detected by the external agent and defeated when the next external agent interaction occurs. If external agent interactions are largely unpredictable to the SDP user, avoiding detection of time fraud on a regular basis will generally require careful monitoring of communication activity, which is likely to necessitate the use of specialized tools for committing fraud.

[0583] Backward Time Detection. In one embodiment the SDP's time service is operable to detect each time it returns a time value that is less than the previous time value it returned. This represents evidence that either the platform time has been adjusted backward as a result of user activity (including automated host activity), or that the time offset has been changed as a result of some interaction with an external agent. Since the time service can be made aware when the latter type of adjustment occurs, unexplained backward motion can be detected unambiguously and appropriate responsive action taken. Note that since time zone information is preferably maintained separately from the SDP instance's indication of absolute time, legitimate changes in time zone will typically not trigger this countermeasure. Storage of the preceding time value is preferably protected against unauthorized modification (e.g., by storing it in tamper-resistant memory and/or by employing obfuscation and/or tamper-resistance techniques such as those described herein). However, finding and modifying the preceding time value generally requires a more sophisticated attack than simply resetting the computer's clock, and thus this counter-

measure can be expected to deter most attackers even if special measures are not taken to protect storage of the preceding time value.

[0584] System Interface Interception. In some embodiments, the SDP instance is operable to detect when the platform's time is explicitly changed, whether through the DATE and TIME operating system commands, through other supported time-setting interfaces (e.g., the system's clock GUI), or through specialized software that invokes low-level interfaces to set the time. If this occurs when the SDP instance is running, the SDP can adjust its time offset accordingly, thus canceling out the attempted time change so that it has no effect on the SDP instance's view of external time. Since this technique detects changes in either direction (i.e., backward and forward), it is generally effective against a wider variety of attempted fraud.

[0585] Moreover, by allowing the external time offset to be adjusted silently when deliberate changes are made, this technique advantageously protects the SDP's integrity while avoiding the need to alert the user about suspected fraud, or to take other intrusive action, when the user may simply be experimenting with the system's clock without fraudulent intent. Nevertheless, it is still typically worthwhile to report the attempt to an external agent, to be used as data in screening for likely fraud.

[0586] If the instance is not running when a deliberate time change occurs, it will typically be necessary to employ additional measures. For example, in one embodiment a special driver can be loaded at system initialization to detect time setting operations and to record an offset adjustment that the instance can fetch when it starts up.

[0587] Instance-Maintained Time. In one embodiment, the SDP is operable to maintain its own idea of time using the platform's timer facility. For example, it might schedule a timer to tick once per second, and detect when two ticks are separated by a much larger (or smaller) interval than one second according to the platform's time service. Such a mechanism would be effective against attempts to set the time explicitly. In addition, in another embodiment a special service is created that does nothing but look for time anomalies. This special service is preferably started at operating system initialization so that stopping the instance server process itself would not interfere.

[0588] Time Anomaly Detection. There may be times during the execution of the SDP instance where an operation that should take a predictable amount of time to complete does not. Such anomalies may indicate attempted time fraud. Accordingly, one countermeasure is to request external time readings relatively frequently during operation of the instance, and to compare the amount of time certain predefined code segments take to execute against predefined thresholds. If an anomaly is detected—for example, if the execution time of a code sequence falls outside of the predefined thresholds—the anomaly can be reported and/or other appropriate defensive action can be taken.

[0589] External Agent Notification. When potential time fraud is detected, it is generally appropriate to

notify an external agent, as discussed previously. This action serves at least two purposes: it notifies the agent of the attempt, and it allows the SDP instance to be resynchronized when the communication with the external agent occurs.

[0590] Time Change Analysis. The countermeasures described above generally address detection of single time-anomaly events. A large anomaly, of course, provides relatively strong evidence that fraud has been attempted, or that gross error has occurred. Similarly, small anomalies are generally benign, as it is reasonable to expect clocks to drift somewhat, and for a user to adjust the machine's clock to compensate. However, an attacker might exploit the fact that a large anomaly can be generated as a sequence of tiny changes. To counter this threat, the SDP instance (or an external agent) can keep track of the history of time anomalies and detect patterns that might result in large changes.

[0591] One such technique is to put a rate limiter on time changes. For example, a user can be prevented from making more than ten changes per day, since even if they are all seemingly benign (e.g., five minutes or less), together they may constitute evidence of attempted fraud and should result in forced external agent notification and/or other defensive action.

[0592] Immediate Verification. Of course, it is possible to require that a time-based SDP action be checked against time obtained directly from the external agent infrastructure, rather than from the platform's time. This is, in effect, an "immediate mode," as it would typically require an online connection to an external agent to supply the time.

[0593] 2.5. Private Storage

[0594] This section describes the use and maintenance of a small amount of memory (for example, 100 bytes) for storing data critical to the SDP's operation. Such private storage can be used to store cryptographic keys and checksums that protect and validate the larger, on-disk database that holds an SDP instance's persistent data. To make the data relatively inaccessible through normal system interfaces, a combination of operating-system-dependent and/or hardware-dependent interfaces which are not ordinarily utilized by application programs—such as SMM on recent Intel processors—can be used.

[0595] Private storage can be used to counter fundamental attacks on an SDP instance (or, in general, on any system that stores transaction information locally). A counter can be maintained in the SDP instance's database. Each time the database is updated, the counter is incremented. Each time the counter is updated, the new counter value is also stored in private storage. If the SDP instance sees that the counter in private storage disagrees with the counter in its own database, it can conclude that there has been some attempt at tampering, and it can enter a tamper response mode.

[0596] It should be appreciated that the data stored in private storage can be more complex than a simple counter. For example, it can be cryptographically derived from the database contents. Moreover, the system can include some fault tolerance and transaction mechanisms to avoid false tampering alarms. The important thing is that each time (or some fraction of times) the database is updated, a value in

private storage is updated also. Private storage can support frequent updates; its purpose is principally to store data for consistency checking, not to hide the data itself. The mechanism may also be used to store information (e.g., cryptographic keys) that should not be disclosed outside the SDP.

[0597] Private storage can be implemented in a variety of ways. In one embodiment the private storage mechanism can use Windows hidden files, or other camouflage techniques, to make it unclear to an attacker what data needs to be copied. Additional implementation techniques are described in the Ginter '900 patent, previously incorporated by reference herein

[0598] 2.6. Driver-Loading Driver

[0599] To avoid tampering with drivers on which the SDP depends, a "driver-loading driver" can be used. This driver can be started at boot-up, and is preferably operable to load other drivers through an interface that is not as easily spoofed, monitored, or intercepted—partly because it is completely undocumented for the attacker, but mostly because it is itself an SDP. This facility provides greater tamper resistance for driver loading, since the load activity will not be as visible and will not require registry changes and so forth to load additional SDP-related drivers. It will also make it easier to distribute new SDP-related driver code transparently, without going through an actual installation process. Moreover, it can be used to deliver updated tamper resistance mechanisms without giving obvious clues like registry changes or updates to driver files.

3. Debugging SDPs

[0600] Once software self-defense techniques have been applied, the resulting program may be more difficult to debug. If debugging an SDP is required, there are several measures that can be applied to make it more practical:

[0601] Since it may not be possible to track down all bugs using the original version of the program, it is generally prudent to retain the ability to regenerate each particular tamper-resistant version of the SDP. Tamper-resistant versions will generally be based on pseudo-random application of software self-defense techniques (possibly in accordance with explicit guidance based on performance and/or security considerations) to an original executable. Therefore, it will be straightforward to reproduce any particular version for support purposes by storing the pseudo-random generator seeds (and any guidance) along with the original executables.

[0602] Provide internal tracing and debugging features. Features for generating (encrypted) trace and debugging messages about the actual semantics of the program's internal operation will generally not be affected by the application of software self-defense techniques. Thus, if they are incorporated in the program, they will typically function as designed, and will support failure diagnosis.

[0603] Rewrite source line number maps. Given information about the equivalence between original instruction locations and their locations in the tamper-resistant version of program, it is relatively straightforward to rewrite the mapping of source code lines to object code locations. Such a mapping permits debuggers to identify the source code associated with an address in a

tamper-resistant program where, for example, an exception was detected. Note that references to source code line numbers made through the C/C++ preprocessor facilities (e.g., the `_LINE_` variable) are not generally affected by these software self-defense techniques, and therefore are still fully meaningful in tamper-resistant versions.

[0604] Rewrite symbol tables. It is also possible to rewrite symbol tables to describe correctly where variables are located in the tamper-resistant versions, thus fully enabling symbolic debugging.

[0605] Provide independent maps. When software self-defense measures result in complex transformations of program behavior, it may not be possible to represent those transformations within the semantics of standard debugging information normally generated for the object file. For example, if a function is replicated in several places, the source line mapping for that function would need to identify several different locations for its generated code, a capability that is often not supported by common symbol table formats. In such cases, it is advantageous for the software self-defense tool to generate an additional database describing the specific results of the software self-defense measures; a tool can be provided to interrogate that database to facilitate debugging. For example, the tool could indicate that a particular instruction address is “replica #3 of code generated for line 783 in file `f_rehash.c`,” that a particular offset in the data section represents a watchdog variable, or so forth.

4. Software Watermarking

[0606] A watermark can be used to deter unauthorized distribution of software by identifying the original owner or licensee of the software. For example, if an attacker anonymously posts a hacked copy of an SDP, a watermark can be used to identify the hacker. A variety of watermarking techniques are described in the Software Watermarking application, which was previously incorporated by reference herein.

[0607] 4.1. Watermark Application

[0608] Watermarks can be applied to an SDP in a variety of ways. In one exemplary embodiment the source code is primed with several assembly language patterns and compiled. The patterns can be found in the executable by performing byte-level pattern matching, and can be replaced (stamped) with a unique watermark on an instance-specific basis. For example, when a user installs the SDP, he may be prompted to connect to a remote server which is operable to decrypt the program and apply a unique watermark.

[0609] In another embodiment, security is improved by moving the watermark writing algorithm from the client to the server. This is achieved by watermarking predefined “patches” at the server and sending them to the client at, e.g., installation. Security is also improved by the fact that the unwatermarked, unencrypted version of the program does not appear at the client.

[0610] This technique is illustrated in FIGS. 35 and 36. As shown in FIG. 36, after a program has been written (3600), the program’s code is primed or marked (3602). Priming can be done at the source code level, although it should be

appreciated that priming could alternatively (or additionally) occur at the object or binary code level. In one embodiment, multiple primers are inserted into the code, each primer being chosen from a predefined set of distinct types (e.g., 10 distinct types). For example, in one embodiment 500 primers of 40 bytes each (on average) are inserted into a 1 megabyte program. It will be recognized that the optimum placement of the primers within the program, as well as the optimum number (and size) of primers to use, will depend on the particular application.

[0611] After compilation (3604), the program executable can be customized (e.g., obfuscated), taking care to preserve the primers (3606). To facilitate this process, the positions of the primers can be made known to the customization tool, so that it can avoid altering them. This can be accomplished in a variety of ways. For example, the customization tool can access the primer reading tool directly (e.g., through DLL calls), and/or the primers can be marked with signal flags indicating that a particular section of code is not to be customized. In another embodiment, a primer reading tool is used to generate a file that provides the locations of the primers, and this file is provided to the customization tool. In yet another embodiment some or all of the primers are simply inserted after customization.

[0612] After customization, in one embodiment each instance of the program is given credentials. This can be accomplished, for example, by computing one or more hash values (using a suitable hashing algorithm, such as MD5 or SHA-1) for one or more program locations, preferably not including the primed regions. As shown in FIG. 35, after the credentialing step, each customized, primed instance of the program can be split into a relatively large part 3504 consisting of most of the executable program, and a smaller part 3506 containing the watermarks or primed locations and/or other security-related code (and an indication of where smaller part 3506 belongs in code 3504) (3608). Code 3504 can contain length preserving holes where the remaining code (e.g., code 3506) would go. In one illustrative embodiment most or all of these holes are filled with NOPs to add stealth. In another embodiment the holes are compressed out, thus effectively disabling code 3504 due to addressing misalignment.

[0613] Code 3504 may be encrypted, packaged, and sent to customers 3512, 3514 and/or distribution centers 3508 (3610, 3612). The appropriate patches 3506 for each customized instance of the program 3504 are stored at a server 3510 in a list 3516 (3614). Thus, when a customer installs program 3504, the appropriate patches 3506 can be supplied by server 3510. For example, as shown in FIG. 35, at installation a customer 3512 contacts server 3510 via a network 3518 such as the Internet. Customer 3512 supplies information regarding the identity (e.g., credentials) of the customized version of the program that the customer received in the installation package. Server 3510 retrieves the corresponding patches 3506, watermarks them (e.g., if they were not already watermarked), compresses and encrypts them (using, e.g., DES or any other suitable encryption technique), and transmits them to the customer. The customer’s system decrypts and decompresses the program and the patches and merges them to form a functioning watermarked program. In a preferred embodiment, the decryption and decompression occur incrementally and concurrently.

[0614] In another embodiment, the patches 3506 can be stored on, e.g., the disk on which program 3504 is stored, and the installation program can apply the appropriate patches. It will be recognized that if patch files are distributed from a server 3510, there will be a tradeoff between security and ease of installation. On the one hand, patch files should be relatively small in order to limit the communication overhead of transmitting them from the server to the client. On the other hand, better security can generally be achieved if the patch files include decoy blocks that contain no watermarks, or if the patch files include random intervals around the watermarks to mask the watermarks' exact locations.

[0615] 4.2. Layered Watermarks

[0616] In one embodiment, multiple watermarks are applied to each instance of the program. As illustrated in FIG. 37A, the container for watermark 3702 includes a pointer or link 3704 operable to indicate the location of another watermark 3706 or the location at which another watermark can be stored.

[0617] One or more unique watermarks can be applied to a program at different stages of the development and distribution process. For example, the software developer might add one or more watermarks to a particular version of the program, these watermarks being the same across all instances of that version of the program. A second watermark can be inserted, for example, when the user installs the program, and can be chosen to be effectively unique for each instance of the program (or for each instance of that version of the program). The watermarks can be located using the version-specific watermark and following a pointer, such as pointer 3704, to the location of the instance-specific watermark. By connecting the watermarks in this manner, the instance specific mark can be relatively small, since it need not contain self-identifying information.

[0618] Although an attacker might be able to detect the instance-specific watermark using a collusive attack, such an attack would typically be unsuccessful in detecting the version-specific watermark since it would be the same across all instances of that version of the program and would not stand out upon comparison of two program instances. Accordingly, this technique provides resistance to collusive attacks, and also provides a level of redundancy that is effective in resisting cropping and patching attacks.

[0619] It will be appreciated that this technique can be readily used to create any suitable number of watermark layers. For example, a program might contain three or more layers of watermarks, the watermarks linked together with pointers in the manner shown in FIG. 37A. From a system perspective, the layers of watermarks create a tree structure, the leaves identifying each instance of the program in an effectively unique fashion. This is illustrated in FIG. 37B, in which node 3752 of tree 3750 represents the original program, nodes 3754 and 3756 represent different versions of the program, nodes 3758 and 3760 represent different instances of those versions, and so forth. Thus it can be seen that each instance of the program can contain a plurality of watermarks—the layers providing collusion resistance, and the multiplicity of watermarks providing resistance to cropping and patching attacks. If an attacker is able to remove one watermark, other watermarks will remain to identify the program.

[0620] 4.3. A Robust Watermark

[0621] To enhance a watermark's robustness, it can be incorporated into actual calculations or operations performed by the program. If an attacker attempts to remove such a watermark, the program may fail to operate correctly. To facilitate adding the watermark to calculations or other program operations, redundancy can be added to the program and/or pre-existing redundancy can be exploited. For example, there are typically a variety of ways to perform any given operation, and thus information can effectively be stored in the choice of which technique is used to perform the operation. Additional examples of redundancy include the jump sequences discussed above with reference to FIGS. 22A through 22D, the manipulation of condition codes at the ends of basic blocks (e.g., using the parity of such variations to encode a watermark), and so forth.

[0622] 4.4. Decoys

[0623] Relatively weak obfuscation, tamper resistance, and/or watermarking techniques can be used to decoy the operation of more effective techniques. For example, decoys can be used to lure the attacker into making program modifications that will break the program or trigger detection of the attacker's activity. For example, having detected a weak watermark or obfuscation technique, the attacker might press forward with eradicating or counteracting it, and not realize that in so doing, he has triggered other, more subtle defensive measures. For example, when removing a weak watermark, the attacker could be lured into making changes to the program binary that would cause the program to fail built-in self-checks, crash or yield incorrect results, and/or trigger communication with an external detection agent.

[0624] Although the foregoing invention has been described in some detail for purposes of clarity of understanding, it will be apparent that certain changes and modifications may be practiced within the scope of the appended claims. It should also be noted that there are many alternative ways of implementing the systems and the methods of the present invention. Accordingly, the present embodiments are to be considered as illustrative and not restrictive, and the invention is not to be limited to the details given herein, but may be modified within the scope and equivalents of the appended claims.

1-78. (canceled)

79. A system for obfuscating computer code, the system comprising:

a memory unit, the memory unit including:

a binary code modification tool;

a program module containing routines for applying predefined obfuscation transformations to computer code; and

a computer program operable to associate the binary code modification tool with the program module to yield an obfuscation program;

a processor for executing the obfuscation program;

wherein the obfuscation program, when executed, is operable to accept computer code as input, to separate the accepted computer code into instructions and data, and to perform transformations on the instructions to pro-

duce an obfuscated computer code, the transformations tending to render the computer code more resistant to reverse engineering, decompilation, or attack.

80. A system as in claim 79, wherein the obfuscation program is further operable to analyze the instructions to determine structure.

81. A system as in claim 80, wherein the obfuscation program is further operable to separate the instructions into basic blocks.

82. A system as in claim 81, wherein the obfuscation program is further operable to identify one or more procedures or functions within one or more basic blocks.

83. A system as in claim 79, wherein the memory unit further includes a compiler for compiling the input computer code.

84. A system as in claim 83, wherein the binary code modification tool is integrated with the compiler.

85. A system as in claim 84, wherein the compiler is controlled by a mode switch, the mode switch allowing the compiler to be switched into or out of a mode in which the binary code modification tool functions to obfuscate code.

86. A system as in claim 84, wherein the compiler applies global software defense techniques at substantially the same time as applying localized software defense techniques.

87. A system as in claim 86, wherein the localized software defense techniques include at least one of techniques for discouraging pattern-matching, techniques for confusing static analysis, techniques for discouraging dynamic analysis, techniques for supplying disinformation, and techniques for applying complex transformations.

88. A system as in claim 83, wherein the compiler generates hints for use by the obfuscation program in obfuscating the computer code.

89. A system as in claim 79, wherein the transformations performed by the obfuscation program on the instructions include at least one of instruction insertion, instruction alteration, and structure alteration.

90. A system as in claim 79, wherein the computer code accepted as input includes one or more of object code and assembly code.

91. As system as in claim 79, wherein the memory unit further includes a source code preprocessor for rearranging at least one of data declarations, and structure or class members.

92. A system as in claim 79, wherein the obfuscation program separates the accepted computer code into instructions and data by, at least in part, performing control flow analysis.

93. A method for obfuscating a computer program, the method including:

linking a binary code modification tool with a program module to form an obfuscation program, wherein the program module contains computer code for performing a plurality of obfuscation transformations;

loading a target program into a memory unit;

separating the target program into instructions and data; and

executing the obfuscation program, the obfuscation program being operable to:

accept at least a portion of the instructions of the target program as input;

perform one or more obfuscation transformations on the instructions; and

write an obfuscated version of the target program using at least part of the transformed instructions to a memory unit as output.

94. A method as in claim 93, wherein the obfuscation program is further operable to analyze the instructions to determine structure.

95. A method as in claim 94, wherein the obfuscation program is further operable to separate the instructions into basic blocks.

96. A method as in claim 95, where in the obfuscation program is further operable to identify one or more procedures or functions within one or more basic blocks.

97. A method as in claim 93 further including:

compiling the target program with a compiler before loading it into the memory unit.

98. A method as in claim 97, wherein the binary code modification tool is integrated with the compiler.

99. A method as in claim 98, further including:

controlling the compiler with a mode switch to allow the compiler to be switched into or out of a mode in which the binary code modification tool functions to obfuscate code.

100. A method as in claim 98, wherein the compiler is operable to apply global software defense techniques at substantially the same time as applying localized software defense techniques.

101. A method as in claim 100, wherein the localized software defense techniques applied by the compiler include at least one of techniques for discouraging pattern-matching, techniques for confusing static analysis, techniques for discouraging dynamic analysis, techniques for supplying disinformation, and techniques for applying complex transformations.

102. A method as in claim 97, wherein the compiler is further operable to generate hints for use by the obfuscation program in obfuscating the target program.

103. A system as in claim 93, wherein the one or more obfuscation transformations performed by the obfuscation program includes at least one of instruction insertion, instruction alteration, and structure alteration.

104. A system as in claim 93, wherein the target program includes one or more of object code and assembly code.

105. A system as in claim 93, further including:

rearranging at least one of data declarations, and structure or class members using a source code preprocessor.

106. A method as in claim 93, further including:

performing control flow analysis to separate the target program into instructions and data.

* * * * *