

(12) STANDARD PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. AU 2013295686 B2

(54) Title
Lock free streaming of executable code data

(51) International Patent Classification(s)
G06F 9/30 (2006.01) **G06F 9/38** (2006.01)

(21) Application No: **2013295686** (22) Date of Filing: **2013.07.26**

(87) WIPO No: **WO14/018812**

(30) Priority Data

(31) Number
13/560,216 (32) Date
2012.07.27 (33) Country
US

(43) Publication Date: **2014.01.30**
(44) Accepted Journal Date: **2018.04.19**

(71) Applicant(s)
Microsoft Technology Licensing, LLC

(72) Inventor(s)
Reierson, Kristofer

(74) Agent / Attorney
Davies Collison Cave Pty Ltd, Level 15 1 Nicholson Street, MELBOURNE, VIC, 3000, AU

(56) Related Art
US 6253309 B1

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(10) International Publication Number

WO 2014/018812 A1

(43) International Publication Date
30 January 2014 (30.01.2014)

(51) International Patent Classification:
G06F 9/30 (2006.01) *G06F 9/38* (2006.01)

(21) International Application Number:
PCT/US2013/052153

(22) International Filing Date:
26 July 2013 (26.07.2013)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
13/560,216 27 July 2012 (27.07.2012) US

(71) Applicant: **MICROSOFT CORPORATION** [US/US];
One Microsoft Way, Redmond, Washington 98052-6399
(US).

(72) Inventor: **REIERSON, Kristofer**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

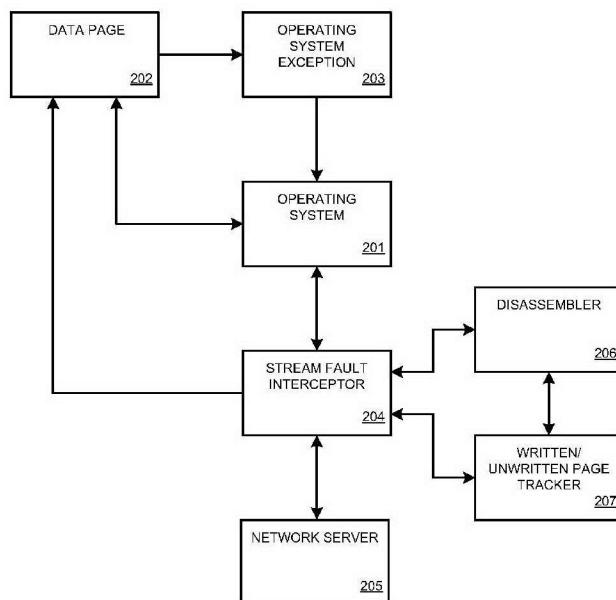
(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

[Continued on next page]

(54) Title: LOCK FREE STREAMING OF EXECUTABLE CODE DATA



(57) Abstract: A disassembler receives instructions and disassembles them into a plurality of separate opcodes. The disassembler creates a table identifying boundaries between each opcode. Each opcode is written to memory in an opcode-by-opcode manner by atomically writing standard blocks of memory. Debug break point opcodes are appended to opcode to create a full block of memory when needed. The block of memory may be thirty-two or sixty-four bits long, for example. Long opcodes may overlap two or more memory blocks. Debug break point opcodes may be appended to a second portion of the long opcode to create a full block of memory. A stream fault interceptor identifies when a requested data page is not available and retrieving the data page.

FIG. 2

2013295686 05 Mar 2018

LOCK FREE STREAMING OF EXECUTABLE CODE DATA

BACKGROUND

[0001] A race condition occurs when commands to read and write data are received at the same time. This can also happen, for example, when a thread is executing a packet of

5 executable data from memory at the same time that a write thread is delivering the data to the memory. The result may be a computer crash, program shutdown, reading/writing errors, or other problems. In one solution, the executing thread may be suspended while waiting for the writing thread to provide the data. However, suspending a thread can cause performance problems for related processes. For example, dozens or hundreds of

10 threads may be running for a process in which only one or a few threads access the data, but all of the threads may have to be suspended to wait for the data delivery every time a stream fault occurs.

[0001A] It is desired to provide a computer-implemented method, a computer system, and a computer-readable storage medium that alleviate one or more difficulties of the prior art, or to at least provide a useful alternative.

SUMMARY

[0001B] In accordance with some embodiments of the present invention, there is provided a computer-implemented method, comprising:

20 receiving instructions comprising a plurality of opcodes, each opcode comprising one or more bytes of data;

disassembling the instructions into separate opcodes;

creating a table identifying the boundaries between each opcode;

appending debug break point opcodes to an opcode to create a full standard block of memory; and

25 writing each opcode to memory in an opcode-by-opcode manner by atomically writing standard blocks of memory.

[0001C] In accordance with some embodiments of the present invention, there is provided a computer system, comprising:

30 a processor;

system memory;

2013295686 05 Mar 2018

one or more computer-readable storage media having stored thereon computer-executable instructions that, when executed by the one or more processors, causes the processors to perform a method for streaming of executable code, the processor operating to:

- 5 receive instructions comprising a plurality of opcodes, each opcode comprising one or more bytes of data;
 - disassembling the instructions into separate opcodes;
 - creating a table identifying the boundaries between each opcode;
 - writing each opcode to memory in an opcode-by-opcode manner by atomically
- 10 writing standard blocks of memory; and
 - append debug break point opcodes to an opcode to create a full standard block of memory.

[0001D] In accordance with some embodiments of the present invention, there is provided a computer-readable storage medium that excludes propagated signals, said computer-readable storage medium storing computer-executable instructions that when executed by at least one processor cause the at least one processor to perform a method for streaming executable code data, the method comprising:

- receiving a fault indicating that a data page is not available;
- 20 requesting instructions including the data page;
- receiving the instructions from a remote source, the instructions comprising a plurality of opcodes, each opcode comprising one or more bytes of data;
 - disassembling the instructions into separate opcodes;
 - creating a table identifying the boundaries between each opcode;
- 25 identifying when a long opcode will overlap two or more standard blocks of memory;
 - writing each opcode to memory in an opcode-by-opcode manner by atomically writing standard blocks of memory;
 - writing a second portion of the long opcode to a second memory block; and
- 30 writing a first portion of the long opcode to a first memory block.

[0002] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended

2013295686 05 Mar 2018

to identify key features or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter.

[0003] Embodiments allow executable code to be delivered in a lock-free manner by disassembling it and ordering the writes to the memory page. This allows the executable code to remain in a consistent state throughout the delivery of streaming data. Because the memory is in a consistent state, other threads may continue to execute while the byte-code is written into memory. This allows for greater scalability and performance.

[0004] Further embodiments enable streaming of executable code data in user mode, without suspending the entire process or using a kernel mode component.

10

DRAWINGS

[0005] Some embodiments of the present invention are hereinafter described, by way of example only, with reference to the accompanying drawings, wherein:

[0006] FIGURES 1A-D illustrate the ordering of the memory writes over time according to one embodiment.

[0007] FIGURE 2 illustrates a system for streaming lock-free executable code data according to one embodiment.

[0008] FIGURE 3 is a flowchart of a process or method for streaming executable code data according to one embodiment.

5 [0009] FIGURE 4 illustrates an example of a suitable computing and networking environment to provide lock free streaming of executable code data.

DETAILED DESCRIPTION

[0010] The lock-free solution described herein takes advantage of two features of a modern processor. First, the processor provides an atomic thirty-two bit or sixty-four bit memory access write. This guarantees that the processor will see all thirty-two bits or sixty-four bits in memory. Second, the processor has a one-byte debug break exception operation code (opcode). This debug break opcode allows the processor to retry reading and executing opcodes from the memory without causing a calling thread to fail.

[0011] In one embodiment, code is disassembled in memory before the instructions are made available to a processor. This disassembly allows for creation of an ordered list of memory writes consisting of a number of separate thirty-two bit or sixty-four bit writes. This disassembly and analysis can be performed either on a client or on a server. Performing this operation on the client saves network transmissions at the cost of increased processor cycles, while performing the operation on the server means larger network transmissions but requires less processor operation at the client node.

[0012] When the memory page is allocated, the entire page is written with one-byte exception codes. A stream fault handler keeps a list of pages and locations that have been written to or that are invalid.

[0013] The opcodes are written to memory using the disassembled thirty-two bit or sixty-four bit blocks starting at the end of the page. Alternately, the assembly tree can be analyzed to identify leaf opcodes, which are opcodes that are called by other opcodes. The leaf opcodes may be identified by a traversal of the opcode execution call graph. These leaf opcodes may be written to memory first.

[0014] For opcodes that are thirty-two bits or smaller (or sixty-four bit or smaller in other embodiments), the entire opcode is written to memory so that the processor will execute the instruction successfully. If the opcode is larger than thirty-two bits (or larger than sixty-four bits in other embodiments), then the “end” of the opcode is written first. By writing the end of the opcode first, if the processor executes the instruction while the opcode is still being written, then the processor will execute the debug break opcode first.

This allows the stream fault handler to finish writing the opcode and so that the processor can retry the opcode. If the thirty-two bit (or sixty-four bit) write straddles more than one opcode, then a debug break opcode is written to allow retry. If the thirty-two bit (or sixty-four bit) write has already been written, then the previous value is written in which case

5 the opcode remains valid.

[0015] By writing the disassembled opcodes spaced by debug breaks, the processor will encounter either an invalid debug breakpoint or valid executable code. If the application executes an invalid debug breakpoint, then a stream fault interceptor as described herein will intercept the exception before it is delivered to the application. The stream fault

10 interceptor waits until the memory is valid and then retries the opcode. The stream fault interceptor can use the table of pages and written opcodes to differentiate between a breakpoint that occurs as a part of stream faulting, which must be retried, or one used by a debugger or the application, which must be passed back to the application. In the case of valid executable code, the application executes code successfully. As a result, the stream

15 fault handler does not need to suspend all of the threads in the process when delivering the results of a stream fault.

[0016] FIGURES 1A-D illustrate the ordering of the memory writes over time according to one embodiment. Program code used in this example comprises the instructions: *0xAB 0xE9 0x00 0xFE 0x70 0x08 . . .* A disassembler breaks this code into the following

20 opcodes:

- (1): *0xAB*
- (2): *0xE9 0x00 0xFE 0x70 0x08*
- (3): *. . .*

The first opcode is one byte long, the second opcode is five bytes long, and additional

25 opcodes of various lengths are identified in a similar manner.

[0017] Memory 101 comprises a number of thirty-two bit blocks 102-104. In other embodiments, the memory blocks may be sixty-four bits or some other size. FIGURE 1A illustrates an original state in which all of memory 101 is filled with one byte debug break point opcodes *0xCC*. If a thread reads instructions from memory 101 at this time, the

30 debug break point opcode will cause the stream handler to attempt to load code to memory 101 and then the thread will retry reading code.

[0018] In FIGURE 1B, the first opcode *0xAB* has been written to thirty-two bit block 102b. The first, one-byte opcode *0xAB* is written into memory 101 with an atomic write that replaces the first *0xCC*. The write operation needs to write thirty-two bits of data.

Because the instruction has been disassembled before writing to memory, the write operation tracks what bytes are being written and knows to fill the remaining twenty-four bits of block 102b with *0xCC* to fill up the entire thirty-two bit block. If the processor executes the first instruction *0xAB*, it will succeed. However, if the processor attempts to 5 execute the second instruction at this time, the debug break point *0xCC* will be read and a stream fault retry mechanism will take over.

[0019] The stream fault handler now needs to write the second opcode (*i.e.*, *0xE9 0x00 0xFE 0X70 0x08*). The stream fault handler knows the previous instruction (*0xAB*) was written and knows how blocks 102b-104b were written. The stream fault handler also 10 knows how long the next opcode is and how it needs to be written to memory 101. In particular, the stream fault handler knows that the second, five-byte opcode will be written to blocks 102b and 103b after opcode *0xAB*. The stream fault handler also determines that the second opcode will not fit into the remaining space in a single thirty-two bit block 102b. Therefore, the second opcode must be divided between memory blocks 102b and 15 103b.

[0020] As illustrated in FIGURE 1C, the second portion of the opcode (*i.e.*, *0x70 0x08*) is written to block 103c – padded with debug break points *0x00* to fill thirty-two bits – before writing the first portion of the opcode. By writing the second portion of the opcode first, if the processor executes the opcode after the *0xAB* block, then the processor will 20 incur a debug fault (*0xCC*) and will not attempt to execute garbage.

[0021] After the second portion of the opcode is written to block 103c, then the first portion (*i.e.*, *0xE9 0x00 0xFE*) can be written to memory 101. As illustrated in FIGURE 1D, this first portion is written with the first opcode *0xAB* in the position that it was previously assigned in block 102d. At this time, the processor will be able to successfully 25 execute both instructions. This process of writing instructions to memory 101 may be applied opcode-by-opcode until the entire page is resident.

[0022] FIGURE 2 illustrates a system for streaming lock-free executable code data according to one embodiment. An application running on operating system 201 attempts to read a data page 202 from memory. When a non-resident page of executable memory is 30 called, this triggers operating system exception 203, which notifies the operating system 201 that this call failed. Operating system 201 sends a read-failure return code to stream fault interceptor 204 before the application sees the fault. Stream fault interceptor 204 looks up the page in its database and determines that it needs to retrieve the page. Stream fault interceptor 204 makes a network call to get the page from network server 205, for

example, and changes the memory access of the page so that it can write the data to data page 202.

[0023] In existing systems, stream fault interceptor 204 would need to suspend the process that requested the data. If the process is not suspended, then another thread could execute data on page 202 as it was being written into memory and the process would crash when it attempted to execute partial data. After the page 202 was written to memory, then the process could be resumed with stream fault interceptor changing the operating system exception to “success” and the application continuing to execute.

[0024] In existing systems, this step of suspending the entire process is necessary because changing access to data page 202 and writing data to it cannot be accomplished in a single step. This means that two threads in the process cannot be executing when you are servicing a stream fault.

[0025] The existing systems are improved herein by adding disassembler 206 and written/unwritten page tracker 207. Disassembler 206 is responsible for taking the returned byte-code and disassembling it. Disassembler 206 returns a list of offsets that contain the boundaries of each opcode. As noted in the example of FIGURE 1, each opcode may be one or more bytes long. Written/unwritten page tracker 207 contains the list of offsets generated by the disassembler 206 as well as a list of those offsets that have been committed.

[0026] The process does not need to be suspended when using disassembler 206 and written/unwritten page tracker 207. Instead, after the requested code is returned to stream fault interceptor 204 from the server 205, the disassembler is called to generate the offset map. The offset map is then stored by the stream fault interceptor 204, which begins writing opcodes to data page 202 in an opcode-by-opcode manner as described above.

The stream fault interceptor 204 consults this offset map to determine the boundary of each opcode. The stream fault interceptor 204 may write the opcodes to data page 202 in any appropriate order starting with either the beginning or the end of the list of opcodes.

[0027] If any contention occurs, such as if the processor executes a partially filled instruction, then the operating system fault mechanism 203 will be triggered. This will cause the stream fault interceptor 204 to be invoked by operating system 201. The stream fault interceptor 204 consults written/unwritten page tracker 207 and either writes the opcode itself to data page 202 or waits until another thread finishes writing the data. Stream fault interceptor 204 then retries the exception.

[0028] FIGURE 3 is a flowchart of a process or method for streaming executable code data according to one embodiment. In step 301, a processor receives instructions comprising a plurality of opcodes. Each of the opcodes may comprise one or more bytes of data. In step 302, a disassembler is used to disassemble the instructions into separate opcodes. In step 303, a table is created that identifies the boundaries between each opcode. The table may be stored in a written/unwritten page tracker, for example. In step 304, each opcode is written to memory in an opcode-by-opcode manner by atomically writing standard blocks of memory. The opcodes may be written to memory by a stream fault interceptor, for example.

5 [0029] In some embodiments, debug break point opcodes may be written to all memory locations to initialize the memory. When the opcodes are written to memory, debug break point opcodes may be appended to the opcode to create a full standard block of memory. The standard block of memory may be thirty-two bits or sixty-four bits long, for example.

10 [0030] When a long opcode will overlap two or more standard blocks of memory, then it may be divided into two portions. The second portion of the long opcode is written to a second memory block, and then the first portion of the long opcode is written to a first memory block. Writing the long opcode in this manner will ensure that a debug break point will be encountered if a memory location is read before the entire opcode is written. Debug break point opcodes may be appended to the second portion of the long opcode to create a full standard block of memory. A previous opcode may be written along with the first portion of the long opcode to the first memory block.

15 [0031] It will be understood that steps 301-305 of the process illustrated in FIGURE 3 may be executed simultaneously and/or sequentially. It will be further understood that each step may be performed in any order and may be performed once or repetitiously.

20 [0032] FIGURE 4 illustrates an example of a suitable computing and networking environment 400 on which the examples of FIGURES 1-3 may be implemented to provide lock free streaming of executable code data. The computing system environment 400 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. The invention is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to: personal computers, server computers, hand-held or laptop devices, tablet devices, multiprocessor systems, microprocessor-based systems, set top

boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

[0033] The invention may be described in the general context of computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, and so forth, which perform particular tasks or implement particular abstract data types. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in local and/or remote computer storage media including memory storage devices.

[0034] With reference to FIGURE 4, an exemplary system for implementing various aspects of the invention may include a general purpose computing device in the form of a computer 400. Components may include, but are not limited to, various hardware components, such as processing unit 401, data storage 402, such as a system memory, and system bus 403 that couples various system components including the data storage 402 to the processing unit 401. The system bus 403 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus also known as Mezzanine bus.

[0035] The computer 400 typically includes a variety of computer-readable media 404. Computer-readable media 404 may be any available media that can be accessed by the computer 400 and includes both volatile and nonvolatile media, and removable and non-removable media, but excludes propagated signals. By way of example, and not limitation, computer-readable media 404 may comprise computer storage media and communication media. Computer storage media includes volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to

store the desired information and which can be accessed by the computer 400.

Communication media typically embodies computer-readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in

5 such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of the any of the above may also be included within the scope of
10 computer-readable media. Computer-readable media may be embodied as a computer program product, such as software stored on computer storage media.

[0036] The data storage or system memory 402 includes computer storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) and random access memory (RAM). Memory 402 or computer readable media 404 may be used to
15 store data pages, opcode boundary lists, opcodes, and the like. A basic input/output system (BIOS), containing the basic routines that help to transfer information between elements within computer 400, such as during start-up, is typically stored in ROM. RAM typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by processing unit 401. By way of example, and not

20 limitation, data storage 402 holds an operating system, application programs, and other program modules and program data. An operating system running on processing unit 402 may support functions such as operating system exception 203, stream fault interceptor 204, disassembler 206, and/or written/unwritten page tracker 207 (FIGURE 2).

[0037] Data storage 402 may also include other removable/non-removable,
25 volatile/nonvolatile computer storage media. By way of example only, data storage 402 may be a hard disk drive that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive that reads from or writes to a removable, nonvolatile magnetic disk, and an optical disk drive that reads from or writes to a removable, nonvolatile optical disk such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary
30 operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The drives and their associated computer storage media, described

above and illustrated in FIGURE 4, provide storage of computer-readable instructions, data structures, program modules and other data for the computer 400.

[0038] A user may enter commands and information through a user interface 405 or other input devices such as a tablet, electronic digitizer, a microphone, keyboard, and/or

5 pointing device, commonly referred to as mouse, trackball or touch pad. Other input devices may include a joystick, game pad, satellite dish, scanner, or the like. Additionally, voice inputs, gesture inputs using hands or fingers, or other natural user interface (NUI) may also be used with the appropriate input devices, such as a microphone, camera, tablet, touch pad, glove, or other sensor. These and other input devices are often connected to the
10 processing unit 401 through a user input interface 405 that is coupled to the system bus 403, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). A monitor 406 or other type of display device is also connected to the system bus 403 via an interface, such as a video interface. The
15 monitor 406 may also be integrated with a touch-screen panel or the like. Note that the monitor and/or touch screen panel can be physically coupled to a housing in which the computing device 400 is incorporated, such as in a tablet-type personal computer. In addition, computers such as the computing device 400 may also include other peripheral output devices such as speakers and printer, which may be connected through an output peripheral interface or the like.

20 [0039] The computer 400 may operate in a networked or cloud-computing environment using logical connections 407 to one or more remote devices, such as a remote computer. The remote computer may be a personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 400. The logical connections depicted in
25 FIGURE 4 include one or more local area networks (LAN) and one or more wide area networks (WAN), but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

[0040] When used in a networked or cloud-computing environment, the computer 400
30 may be connected to a public or private network through a network interface or adapter 407. Network interface 407 may provide a connection to a remote device, such as a network server 205 (FIGURE 2). In some embodiments, a modem or other means for establishing communications over the network. The modem, which may be internal or external, may be connected to the system bus 403 via the network interface 407 or other

2013295686 05 Mar 2018

appropriate mechanism. A wireless networking component such as comprising an interface and antenna may be coupled through a suitable device such as an access point or peer computer to a network. In a networked environment, program modules depicted relative to the computer 400, or portions thereof, may be stored in the remote memory storage device. It may be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

5 [0041] Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined 10 in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims.

15 [0042] Throughout this specification and claims which follow, unless the context requires otherwise, the word "comprise", and variations such as "comprises" and "comprising", will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps.

20 [0043] The reference in this specification to any prior publication (or information derived from it), or to any matter which is known, is not, and should not be taken as an acknowledgment or admission or any form of suggestion that that prior publication (or information derived from it) or known matter forms part of the common general knowledge in the field of endeavour to which this specification relates.

THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1. A computer-implemented method, comprising:

receiving instructions comprising a plurality of opcodes, each opcode comprising one or more bytes of data;

5 disassembling the instructions into separate opcodes;

creating a table identifying the boundaries between each opcode;

appending debug break point opcodes to an opcode to create a full standard block of memory; and

writing each opcode to memory in an opcode-by-opcode manner by atomically

10 writing standard blocks of memory.

2. The computer-implemented method of claim 1, further comprising: writing debug break point opcodes to all memory locations to initialize the memory.

15 3. The computer-implemented method of claim 1, wherein the standard block of memory is thirty-two bits or sixty-four bits long.

4. The computer-implemented method of claim 1, further comprising: identifying when a long opcode will overlap two or more standard blocks of memory; writing a

20 second portion of the long opcode to a second memory block; and writing a first portion of the long opcode to a first memory block.

5. The computer-implemented method of claim 4, further comprising: appending debug break point opcodes to the second portion of the long opcode to create a full

25 standard block of memory.

6. The computer-implemented method of claim 5, further comprising: writing a previous opcode and the first portion of the long opcode to the first memory block.

30 7. A computer system, comprising:

a processor;

system memory;

2013295686 05 Mar 2018

one or more computer-readable storage media having stored thereon computer-executable instructions that, when executed by the one or more processors, causes the processors to perform a method for streaming of executable code, the processor operating to:

- 5 receive instructions comprising a plurality of opcodes, each opcode comprising one or more bytes of data;
 - disassembling the instructions into separate opcodes;
 - creating a table identifying the boundaries between each opcode;
 - writing each opcode to memory in an opcode-by-opcode manner by atomically
- 10 writing standard blocks of memory; and
 - append debug break point opcodes to an opcode to create a full standard block of memory.

8. The computer system of claim 7, the processor further operating to: writing debug
15 break point opcodes to all memory locations to initialize the memory.

9. The computer system of claim 7, wherein the standard block of memory is thirty-two bits or sixty-four bits long.

20 10. The computer system of claim 7, the processor further operating to: identify when a long opcode will overlap two or more standard blocks of memory; write a second portion of the long opcode to a second memory block; and write a first portion of the long opcode to a first memory block.

25 11. The computer system of claim 10, the processor further operating to: append debug break point opcodes to the second portion of the long opcode in to create a full standard block of memory.

30 12. The computer system of claim 10, the processor further operating to: write a previous opcode and the first portion of the long opcode to the first memory block.

13. The computer system of claim 7, further comprising: a stream fault interceptor operating to identify when a requested data page is not available; and retrieving the data page from a remote source.

14. The computer system of claim 7, further comprising: a disassembler operating to disassemble the instructions into separate opcodes.

5 15. The computer system of claim 7, further comprising: a page tracker maintaining the table that identifies the boundaries between each opcode.

16. A computer-readable storage medium that excludes propagated signals, said computer-readable storage medium storing computer-executable instructions that when 10 executed by at least one processor cause the at least one processor to perform a method for streaming executable code data, the method comprising:

receiving a fault indicating that a data page is not available;

requesting instructions including the data page;

receiving the instructions from a remote source, the instructions comprising a

15 plurality of opcodes, each opcode comprising one or more bytes of data;

disassembling the instructions into separate opcodes;

creating a table identifying the boundaries between each opcode;

identifying when a long opcode will overlap two or more standard blocks of 20 memory;

writing each opcode to memory in an opcode-by-opcode manner by atomically writing standard blocks of memory;

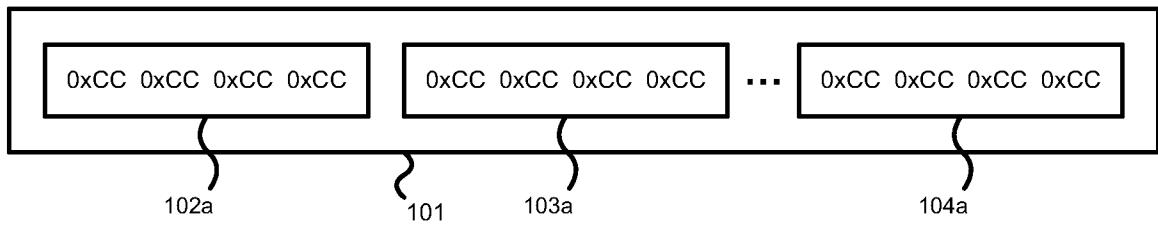
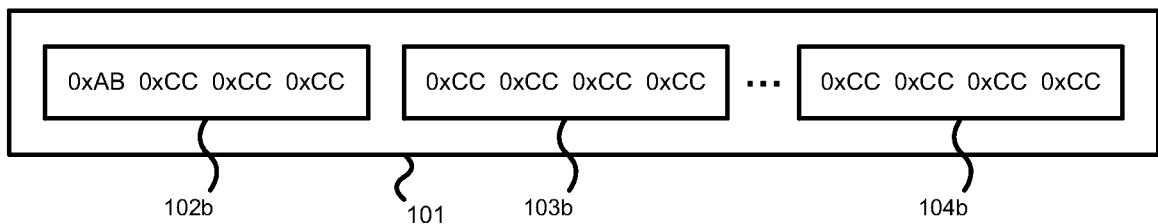
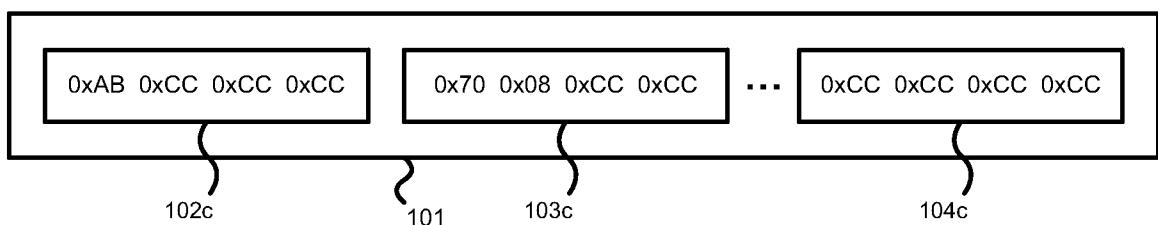
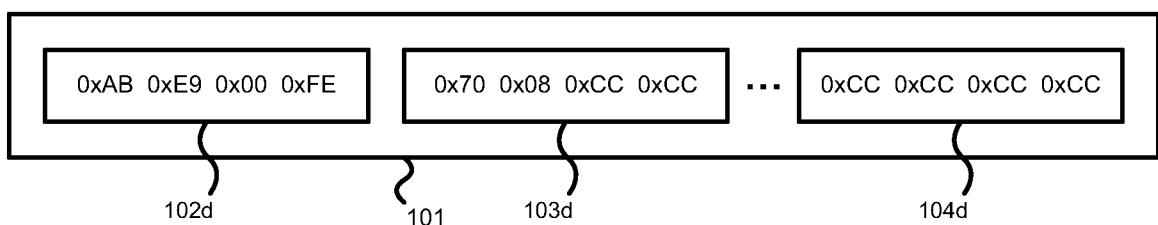
writing a second portion of the long opcode to a second memory block; and

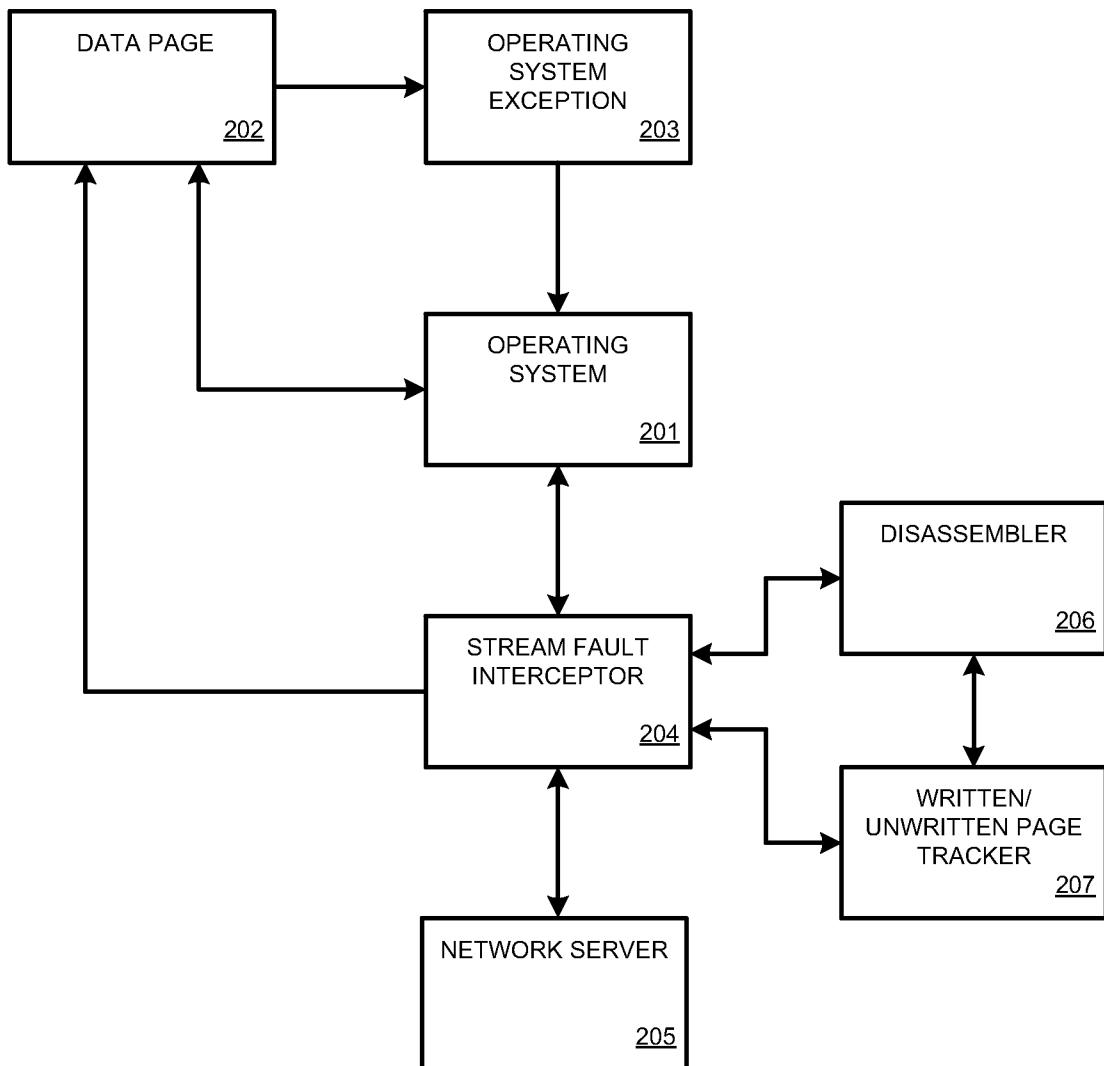
writing a first portion of the long opcode to a first memory block.

25 17. The computer-readable storage medium of claim 16, wherein the method further comprises: appending debug break point opcodes to an opcode to create a full standard block of memory.

18. The computer-readable storage medium of claim 16, wherein the standard block of 30 memory is thirty-two bits or sixty-four bits long.

1/3

**FIG. 1A****FIG. 1B****FIG. 1C****FIG. 1D**

**FIG. 2**

3/3

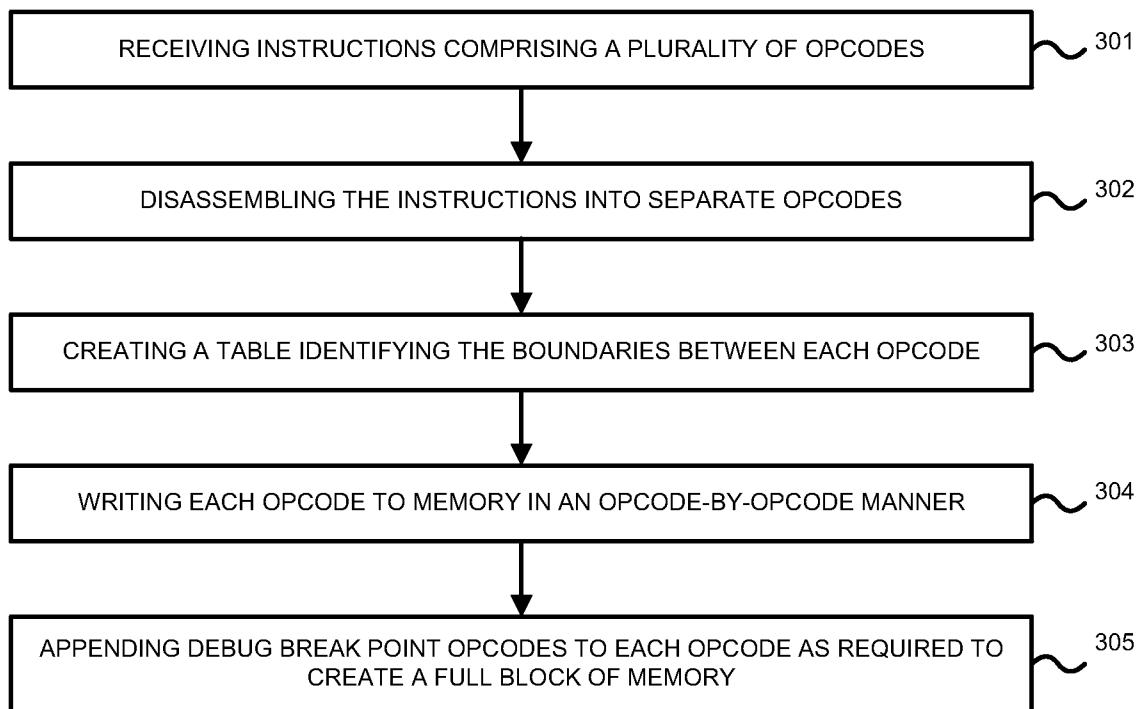


FIG. 3

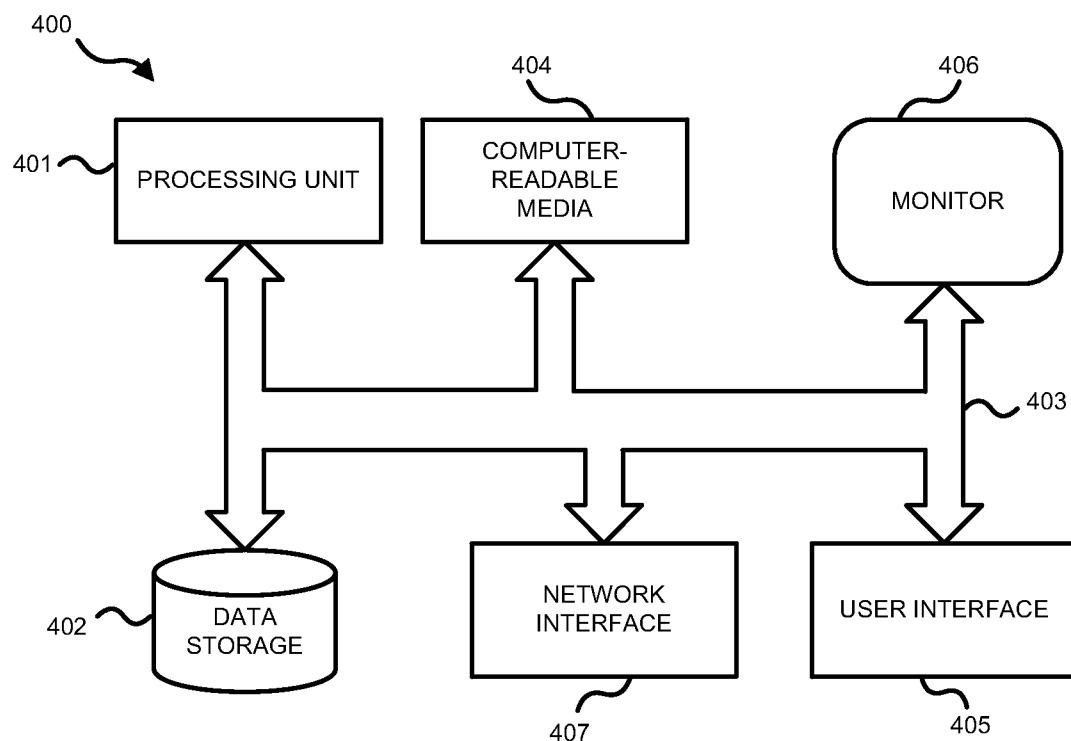


FIG. 4