

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
24 March 2011 (24.03.2011)

(10) International Publication Number  
**WO 2011/035300 A2**

(51) International Patent Classification:  
*G06F 9/50* (2006.01)

Victory Memorial Dr., Minneapolis, Minnesota 55412 (US).

(21) International Application Number:  
PCT/US2010/049650

(74) Agents: **WALLACE, JR., Michael J.** et al.; Siemens Corporation- Intellectual Property Dept., 170 Wood Avenue South, Iselin, New Jersey 08830 (US).

(22) International Filing Date:  
21 September 2010 (21.09.2010)

(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, 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, KM, 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, PE, PG, PH, PL, PT, RO, RS, RU, 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.

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
12/563,614 21 September 2009 (21.09.2009) US

(71) Applicant (for all designated States except US):  
**SIEMENS PRODUCT LIFECYCLE MANAGEMENT SOFTWARE INC.** [US/US]; 5800 Granite Parkway, Suite 600, Plano, Texas 75024 (US).

(72) Inventors; and

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

(75) Inventors/Applicants (for US only): **WHELAN, John Staehle** [US/US]; 2121 8th Street, White Bear Lake, Minnesota 55110 (US). **LUDWIG, Mark** [US/US]; 2800

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR TRANSFER OF DATA STRUCTURES WITH TRANSFER STATE AWARENESS

(57) Abstract: A method for data transfer in a data processing system, and corresponding system and machine-readable medium. One method includes receiving by the data processing system a request for a data structure from a calling process, and splitting the data structure into a plurality of substructures by the data processing system. That method includes transferring the plurality of substructures to the calling process by the data processing system, wherein at least two of the substructures are transferred in parallel, and maintaining a transfer state for each substructure in the data processing system.




Figure 2

200

WO 2011/035300 A2

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, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Published:**

— *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

SYSTEM AND METHOD FOR TRANSFER OF DATA  
STRUCTURES WITH TRANSFER STATE AWARENESS

TECHNICAL FIELD

[0001] The present disclosure is directed, in general, to data transfer and, more specifically, to efficient transfer of large data structures.

BACKGROUND OF THE DISCLOSURE

[0002] Transferring large data structures or large blocks of structured data can be time consuming and act as a bottleneck to efficient data processing.

SUMMARY OF THE DISCLOSURE

[0003] A method for data transfer in a data processing system, and corresponding system and machine-readable medium. One method includes receiving by the data processing system a request for a data structure from a calling process, and splitting the data structure into a plurality of substructures by the data processing system. That method includes transferring the plurality of substructures to the calling process by the data processing system, wherein at least two of the substructures are transferred in parallel, and maintaining a transfer state for each substructure in the data processing system.

[0004] The foregoing has outlined rather broadly the features and technical advantages of the present disclosure so that those skilled in the art may better understand the detailed description that follows. Additional features and advantages of the disclosure will be described hereinafter that form the subject of the claims. Those skilled in the art will appreciate that they may readily use the conception and the specific embodiment disclosed as a basis for modifying or designing other structures for carrying out the same purposes of the present disclosure. Those skilled

in the art will also realize that such equivalent constructions do not depart from the spirit and scope of the disclosure in its broadest form.

[0005] Before undertaking the DETAILED DESCRIPTION below, it may be advantageous to set forth definitions of certain words or phrases used throughout this patent document: the terms “include” and “comprise,” as well as derivatives thereof, mean inclusion without limitation; the term “or” is inclusive, meaning and/or; the phrases “associated with” and “associated therewith,” as well as derivatives thereof, may mean to include, be included within, interconnect with, contain, be contained within, connect to or with, couple to or with, be communicable with, cooperate with, interleave, juxtapose, be proximate to, be bound to or with, have, have a property of, or the like; and the term “controller” means any device, system or part thereof that controls at least one operation, whether such a device is implemented in hardware, firmware, software or some combination of at least two of the same. It should be noted that the functionality associated with any particular controller may be centralized or distributed, whether locally or remotely. Definitions for certain words and phrases are provided throughout this patent document, and those of ordinary skill in the art will understand that such definitions apply in many, if not most, instances to prior as well as future uses of such defined words and phrases. While some terms may include a wide variety of embodiments, the appended claims may expressly limit these terms to specific embodiments.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0006] For a more complete understanding of the present disclosure, and the advantages thereof, reference is now made to the following descriptions taken in conjunction with the accompanying drawings, wherein like numbers designate like objects, and in which:

[0007] Figure 1 depicts a block diagram of a data processing system in which an embodiment can be implemented;

[0008] Figure 2 illustrates a data transfer process as performed by a data processing system in accordance with disclosed embodiments;

[0009] Figure 3 illustrates an example of data structure splitting in accordance with disclosed embodiments; and

[0010] Figure 4 depicts a process in accordance with disclosed embodiments.

#### DETAILED DESCRIPTION

[0011] FIGURES 1 through 4, discussed below, and the various embodiments used to describe the principles of the present disclosure in this patent document are by way of illustration only and should not be construed in any way to limit the scope of the disclosure. Those skilled in the art will understand that the principles of the present disclosure may be implemented in any suitably arranged device. The numerous innovative teachings of the present application will be described with reference to exemplary non-limiting embodiments.

[0012] Transferring large data structures or large blocks of structured data between applications or systems is a problem that is currently solved in linear or greater time relative to the size of the structure being transferred. As data structures become larger, these time limits become bottlenecks to timely sharing of data. Additionally, the consumption of resources for transferring a structure, including CPU cycles and memory, grows in a linear or greater manner relative to the size of the structure being transferred, and is restricted to operating on a single computer. The restriction of resources to a single computer produces a limitation to the ability to support the processing of arbitrary large data structures due to upper limits to the size of a computer.

[0013] Disclosed embodiments provide a system and method for transferring data between differing processes or applications executing on a single system or on multiple interconnected data processing systems, each having one or more processors.

[0014] Figure 1 depicts a block diagram of a data processing system in which an embodiment can be implemented, as any of the systems operating on data structures

as described herein. The data processing system depicted includes one or more processors such as a processor 102 connected to a level two cache/bridge 104, which is connected in turn to a local system bus 106. Note that although only one processor 102 is shown in this diagram for simplicity, various embodiments explicitly include multiple processors. Local system bus 106 may be, for example, a peripheral component interconnect (PCI) architecture bus. Also connected to local system bus in the depicted example are a main memory 108 and a graphics adapter 110. The graphics adapter 110 may be connected to display 111.

**[0015]** Other peripherals, such as local area network (LAN) / Wide Area Network / Wireless (*e.g.* WiFi) adapter 112, may also be connected to local system bus 106. Expansion bus interface 114 connects local system bus 106 to input/output (I/O) bus 116. I/O bus 116 is connected to keyboard/mouse adapter 118, disk controller 120, and I/O adapter 122. Disk controller 120 can be connected to a storage 126, which can be any suitable machine usable or machine readable storage medium, including but not limited to nonvolatile, hard-coded type mediums such as read only memories (ROMs) or erasable, electrically programmable read only memories (EEPROMs), magnetic tape storage, and user-recordable type mediums such as floppy disks, hard disk drives and compact disk read only memories (CD-ROMs) or digital versatile disks (DVDs), and other known optical, electrical, or magnetic storage devices.

**[0016]** Also connected to I/O bus 116 in the example shown is audio adapter 124, to which speakers (not shown) may be connected for playing sounds. Keyboard/mouse adapter 118 provides a connection for a pointing device (not shown), such as a mouse, trackball, trackpointer, etc.

**[0017]** Those of ordinary skill in the art will appreciate that the hardware depicted in Figure 1 may vary for particular implementations. For example, other peripheral devices, such as an optical disk drive and the like, also may be used in addition or in place of the hardware depicted. The depicted example is provided for the purpose of explanation only and is not meant to imply architectural limitations with respect to the present disclosure.

**[0018]** A data processing system in accordance with an embodiment of the present disclosure includes an operating system employing a graphical user interface. The

operating system permits multiple display windows to be presented in the graphical user interface simultaneously, with each display window providing an interface to a different application or to a different instance of the same application. A cursor in the graphical user interface may be manipulated by a user through the pointing device. The position of the cursor may be changed and/or an event, such as clicking a mouse button, generated to actuate a desired response.

**[0019]** One of various commercial operating systems, such as a version of Microsoft Windows™, a product of Microsoft Corporation located in Redmond, Wash. may be employed if suitably modified. The operating system is modified or created in accordance with the present disclosure as described.

**[0020]** LAN/ WAN/Wireless adapter 112 can be connected to a network 130 (not a part of data processing system 100), which can be any public or private data processing system network or combination of networks, as known to those of skill in the art, including the Internet. Data processing system 100 can communicate over network 130 with server system 140, which is also not part of data processing system 100, but can be implemented, for example, as a separate data processing system 100. Various embodiments described herein can include the data transfers between multiple data processing systems over a network 130, each of the data processing systems implemented as a data processing system 100.

**[0021]** Organizations with large quantities of structured data wish to share data both internally and externally. Often, the sharing of data needs to be accomplished in a specific amount of time. The ability to share data faster provides a significant technical advantage and significant business advantages to an organization.

**[0022]** Various embodiments include a system and method for splitting individual transfers of structured data into multiple sub-requests that can run independently in parallel, in order to achieve support of transfer of arbitrary size and in sub-linear time. This approach provides significant improvements to systems such as enterprise information management systems, by enabling the system to remember what data has been transferred (“transfer state awareness”), and to transfer portions of data containing only the data that has not been sent.

[0023] The disclosed embodiments include adding additional steps to a data transfer. These include analyzing the structure being transferred, and recursively subdividing the structure into transfers of bounded-size subsets of the structure for which non-dependent sub-structures can be transferred in parallel. This approach allows the work to be horizontally scaled across an arbitrary number of processors, and to have transfers completed in logarithmic time (given sufficient processors).

[0024] Figure 2 illustrates a data transfer process as performed by a data processing system 200 in accordance with disclosed embodiments. As shown here, data processing system 200 responds to a data request from caller 205 by invoking multiple instances of a transfer engine 210 and invoking a transfer splitter 215. Data processing system 200 also maintains a transfer state awareness for the portions of the data being transferred. Data processing system 200 can be implemented, for example, as a data processing system 100.

[0025] This figure illustrates the recursive nature of the disclosed process. The caller 205 requests data. TransferEngine 210 calls the TransferSplitter 215 at 220, which splits the requested data structure into multiple substructures. Substructures are identified in the split actions, and are sent for processing as separate calls to other instances of TransferEngine 205 at 225, where there may be multiple calls 225. Upon completion of the calls for each substructure, as indicated at 230, the parent structure is then processed in a recursive manner at 235.

[0026] The actual transfer of the various substructures can be performed in a conventional manner, with the addition of maintaining transfer state awareness.

[0027] Figure 3 illustrates an example of data structure splitting in accordance with disclosed embodiments. Structure 305 is the starting element, and is shown with a three subelements. In this particular case, the structure 305 is being split into four substructures 310, each with a size of three. Each of these substructures can be recursively broken down further into substructures 315.

[0028] Each of the substructures can be assigned to a specific TransferEngine instance that processes a specific subset of the structure. In this example, there are three TransferEngine instances designated A, B, and C used to transfer the structure. Each of the corresponding structures and substructures in Figure 3 is illustrated with



an A, B, or C label showing the TransferEngine instance to which it is assigned. Since the TransferEngine instances transfer the substructures in parallel, each set of three A, B, and C substructures can be transferred at roughly the same time, rather than in a serial manner if the entire structure were transferred at once by a single transfer process as in conventional systems.

**[0029]** Adding more TransferEngine instances will decrease the amount of time needed to transfer a structure, with a lower limit being related to the depth of a structure.

**[0030]** Following is pseudo-code of a process that can be used for structure splitting in accordance with some embodiments:

```

Transfer.process (TransferRequest x) {
    // holds onto the request until the time is should run.
    SchedulingService.process(x);

    // splits off subparts of the transfer to other transfers
    Integer transactionSize = 1000;
    SplittingService.process(x, transactionSize);

    // The "normal" steps..
    ExportContents ec = ExportService.process(x);
    MappedContents mc = MappingService.process(ec, x);
    ImportContents ic = ImportService.process(mc, x);
    ConfirmExportService.process(ic, x);
}

SplittingService.process (TransferRequest x, Integer
transactionSize) {
    PDMObject rootObj = getRootObject(x);

    // traverse down the structure until you reach a size of
    transactionSize
    PDMObject leftObjs[] = { rootObj};
    Integer totalSize = 1;
    Do {
        leftObjs = BackEndSystem.setExpand( leftObjs,
"RelationshipName" );
        totalSize = totalSize + leftObjs.size;
    }
}

```

```

    } While ( totalSize < transactionSize )

    For Each (obj: leftObjs) {
        // generate request for child object
        TransferRequest y = generateTransferRequest(x,
obj);

        // RECURSION HERE! (async call with callback)
        TransferBPPEL.process(y);
    }

    // wait for all the child transfers to complete.
    Wait(leftObjs);

    Return;
}

```

**[0031]** The exemplary pseudo-code shows one sample process that can be used for splitting data structures as described herein. The substructures identified by the splitter are transferred in parallel. The recursive nature of the process allows for arbitrary large structures to be processed. The parent structure(s) await completion of the substructures, thus allowing for failures in the sub-transactions to not result in missed data by the parent transfer(s).

**[0032]** Because subsequent steps are being processed as sub-structures, even though the steps themselves are linear in cost, because they are being run in parallel, they can be processed in sub-linear time. The recursion approach allows independent sub-structures (sub-structures that are not ancestor/descendent related) can be run in parallel, while dependent (parent) structures are always run after the lower level structures are completed. This allows the system to avoid redundant traversals and exports for systems that maintain transfer-state awareness (*i.e.*, the system remembers that it already transferred some element's subelement, and does not redundantly re-transfer).

**[0033]** This algorithm will scale to multiple processors, and have an lower wall-clock time limit of  $WallClockTime = \text{Log}X(N)+C$ , where  $X$  is related to the average number children per node,  $N$  is the number of nodes in the structure, and  $C$  is a constant overhead.

[0034] Figure 4 depicts a process in accordance with disclosed embodiments.

[0035] The data processing system receives a request (or call) for a data structure in a transfer engine and from a calling process (step 405).

[0036] The data processing system splits the data structure into a plurality of substructures (step 410). Each (child) substructure can have a dependent parent structure or substructure. The splitting process can be performed recursively on each substructure to a selected granularity that can be specified either in number or size of the smallest substructures, or according to the depth of the substructure tree. This step can include calling a transfer splitter engine by the transfer engine.

[0037] The data processing system transfers the substructures to the calling process, wherein at least two of the substructures are transferred in parallel (step 415). Preferably, each substructure transferred in parallel is transferred by a separate processor or processor core. This step can include transferring each substructure using an independent transfer engine instance.

[0038] The data processing system maintains a record of a transfer state of each of the substructures until the transfer of all substructures is complete (step 420). The transfer state can include states such as transferred, transfer failed, and untransferred. Subsequent substructure transfers can include any substructures having a state of transfer failed or untransferred, but preferably does not include any substructures having a state indicating they have already been transferred. The maintaining step is necessarily performed at the same time as the transferring step.

[0039] The calling process optionally re-assembles the data structure from the substructures (step 425).

[0040] Those skilled in the art will recognize that, for simplicity and clarity, the full structure and operation of all data processing systems suitable for use with the present disclosure is not being depicted or described herein. Instead, only so much of a data processing system as is unique to the present disclosure or necessary for an understanding of the present disclosure is depicted and described. The remainder of the construction and operation of data processing system 100 may conform to any of the various current implementations and practices known in the art.

[0041] It is important to note that while the disclosure includes a description in the context of a fully functional system, those skilled in the art will appreciate that at least portions of the mechanism of the present disclosure are capable of being distributed in the form of a instructions contained within a machine-usable, computer-usable, or computer-readable medium in any of a variety of forms, and that the present disclosure applies equally regardless of the particular type of instruction or signal bearing medium or storage medium utilized to actually carry out the distribution. Examples of machine usable/readable or computer usable/readable mediums include: nonvolatile, hard-coded type mediums such as read only memories (ROMs) or erasable, electrically programmable read only memories (EEPROMs), and user-recordable type mediums such as floppy disks, hard disk drives and compact disk read only memories (CD-ROMs) or digital versatile disks (DVDs).

[0042] Although an exemplary embodiment of the present disclosure has been described in detail, those skilled in the art will understand that various changes, substitutions, variations, and improvements disclosed herein may be made without departing from the spirit and scope of the disclosure in its broadest form.

[0043] None of the description in the present application should be read as implying that any particular element, step, or function is an essential element which must be included in the claim scope: the scope of patented subject matter is defined only by the allowed claims. Moreover, none of these claims are intended to invoke paragraph six of 35 USC §112 unless the exact words "means for" are followed by a participle.

## WHAT IS CLAIMED IS:

1. A method for data transfer in a data processing system, comprising:  
receiving by the data processing system a request for a data structure from a calling process;  
splitting the data structure into a plurality of substructures by the data processing system;  
transferring the plurality of substructures to the calling process by the data processing system, wherein at least two of the substructures are transferred in parallel; and  
maintaining a transfer state for each substructure in the data processing system.
2. The method of claim 1, wherein the maintaining step is performed concurrently with the transferring step.
3. The method of claim 1, wherein the transferring step uses an independent instance of a transfer engine for each substructure.
4. The method of claim 1, wherein the splitting is performed recursively to produce substructures of a selected size.
5. The method of claim 1, wherein substructures that are transferred in parallel are each transferred using a different processor of the data processing system.
6. The method of claim 1, wherein the calling process reassembles the data structure from the substructures.
7. The method of claim 1, wherein the state for each substructure indicates whether the substructure has been successfully transferred.
8. A data processing system comprising a plurality of processors and an accessible memory, the data processing system particularly configured to perform the steps of:

receiving a request for a data structure from a calling process;  
splitting the data structure into a plurality of substructures;  
transferring the plurality of substructures to the calling process, wherein at  
least two of the substructures are transferred in parallel; and  
maintaining a transfer state for each substructure.

9. The data processing system of claim 8, wherein the maintaining step is performed concurrently with the transferring step.
10. The data processing system of claim 8, wherein the transferring step uses an independent instance of a transfer engine for each substructure.
11. The data processing system of claim 8, wherein the splitting is performed recursively to produce substructures of a selected size.
12. The data processing system of claim 8, wherein substructures that are transferred in parallel are each transferred using a different processor of the data processing system.
13. The data processing system of claim 8, wherein the calling process reassembles the data structure from the substructures.
14. The data processing system of claim 8, wherein the state for each substructure indicates whether the substructure has been successfully transferred.
15. A tangible machine-readable medium encoded with computer-executable instructions that, when executed, cause a data processing system to perform the steps of:  
receiving a request for a data structure from a calling process;  
splitting the data structure into a plurality of substructures;  
transferring the plurality of substructures to the calling process, wherein at  
least two of the substructures are transferred in parallel; and  
maintaining a transfer state for each substructure.

16. The machine-readable medium of claim 15, wherein the maintaining step is performed concurrently with the transferring step.
17. The machine-readable medium of claim 15, wherein the transferring step uses an independent instance of a transfer engine for each substructure.
18. The machine-readable medium of claim 15, wherein the splitting is performed recursively to produce substructures of a selected size.
19. The machine-readable medium of claim 15, wherein substructures that are transferred in parallel are each transferred using a different processor of the data processing system.
20. The machine-readable medium of claim 15, wherein the state for each substructure indicates whether the substructure has been successfully transferred.

100

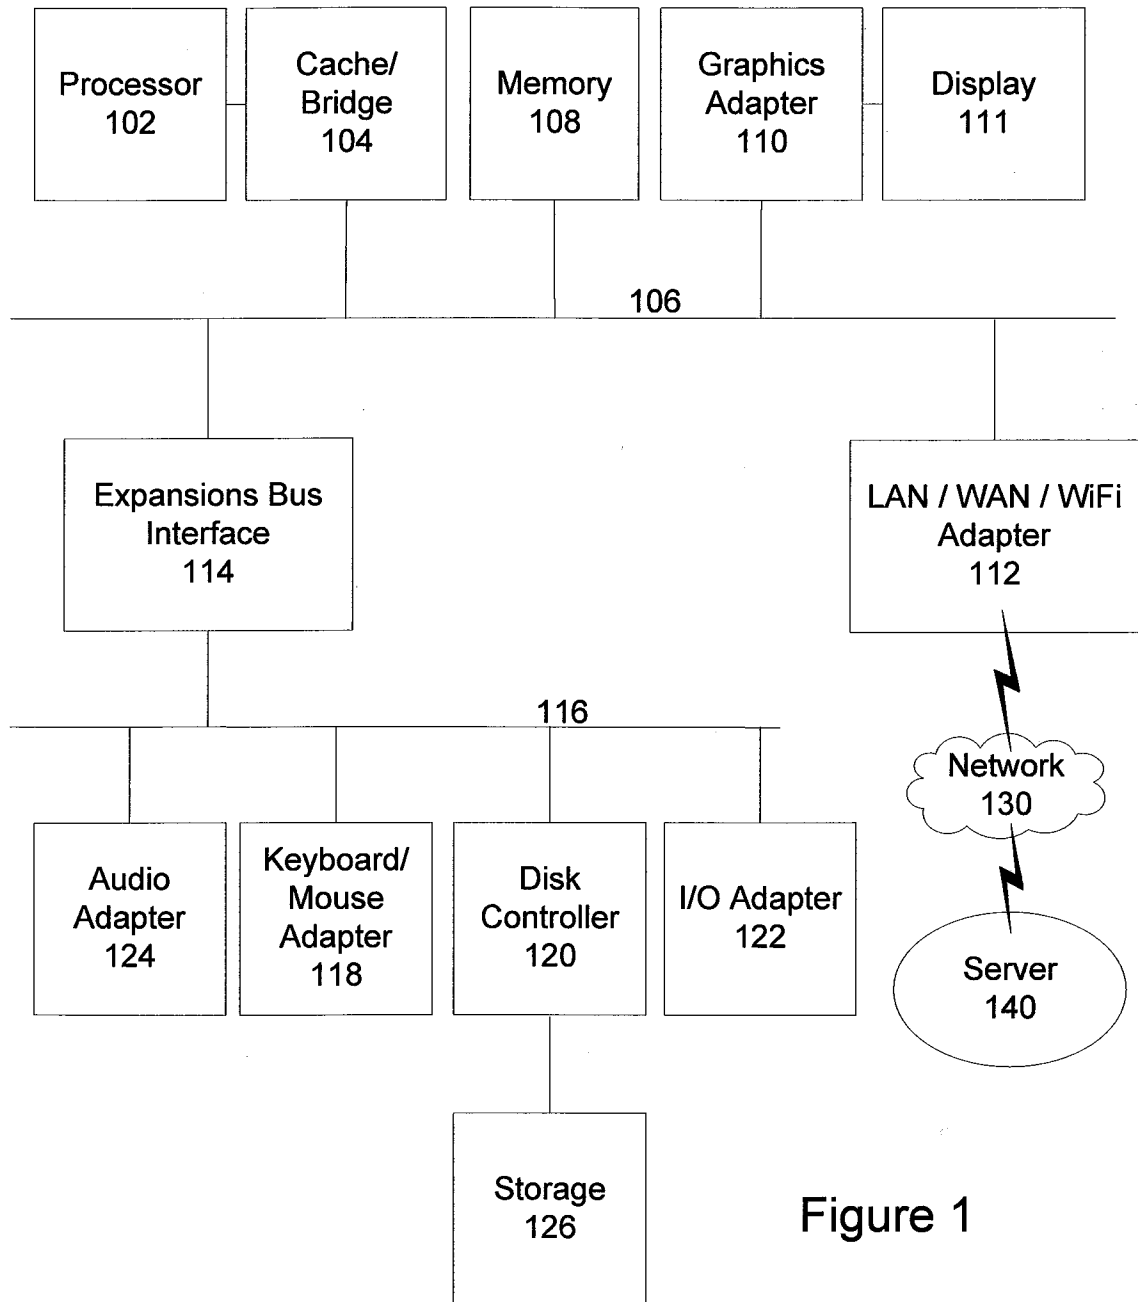


Figure 1



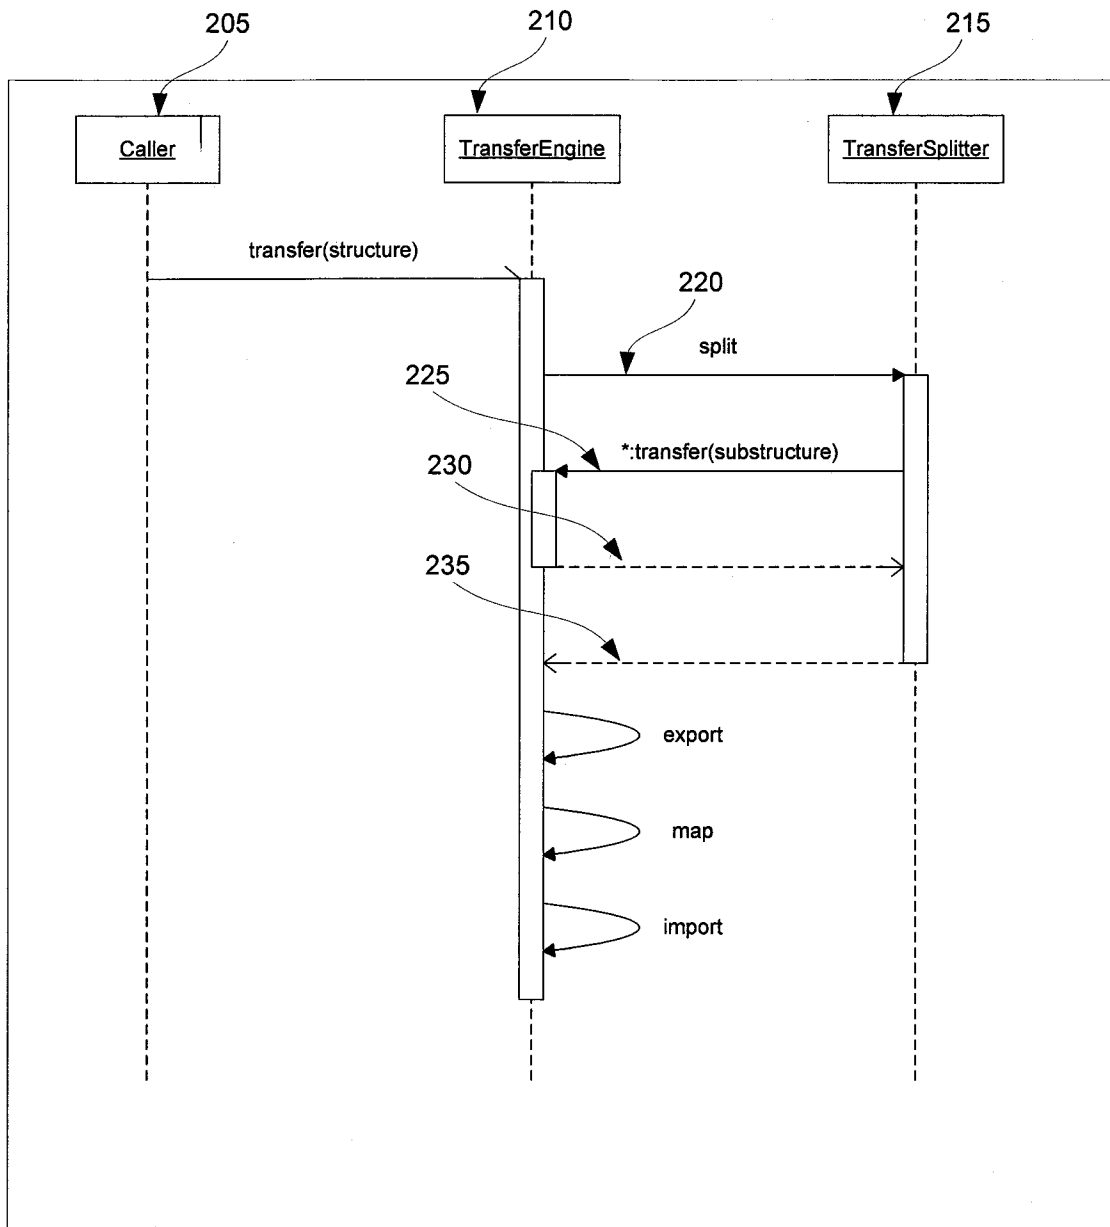


Figure 2

200




Figure 3

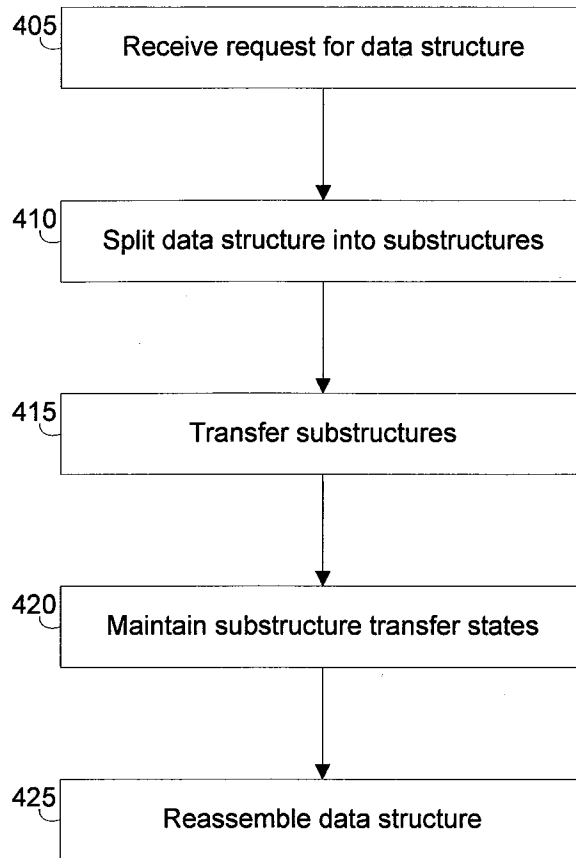


Figure 4