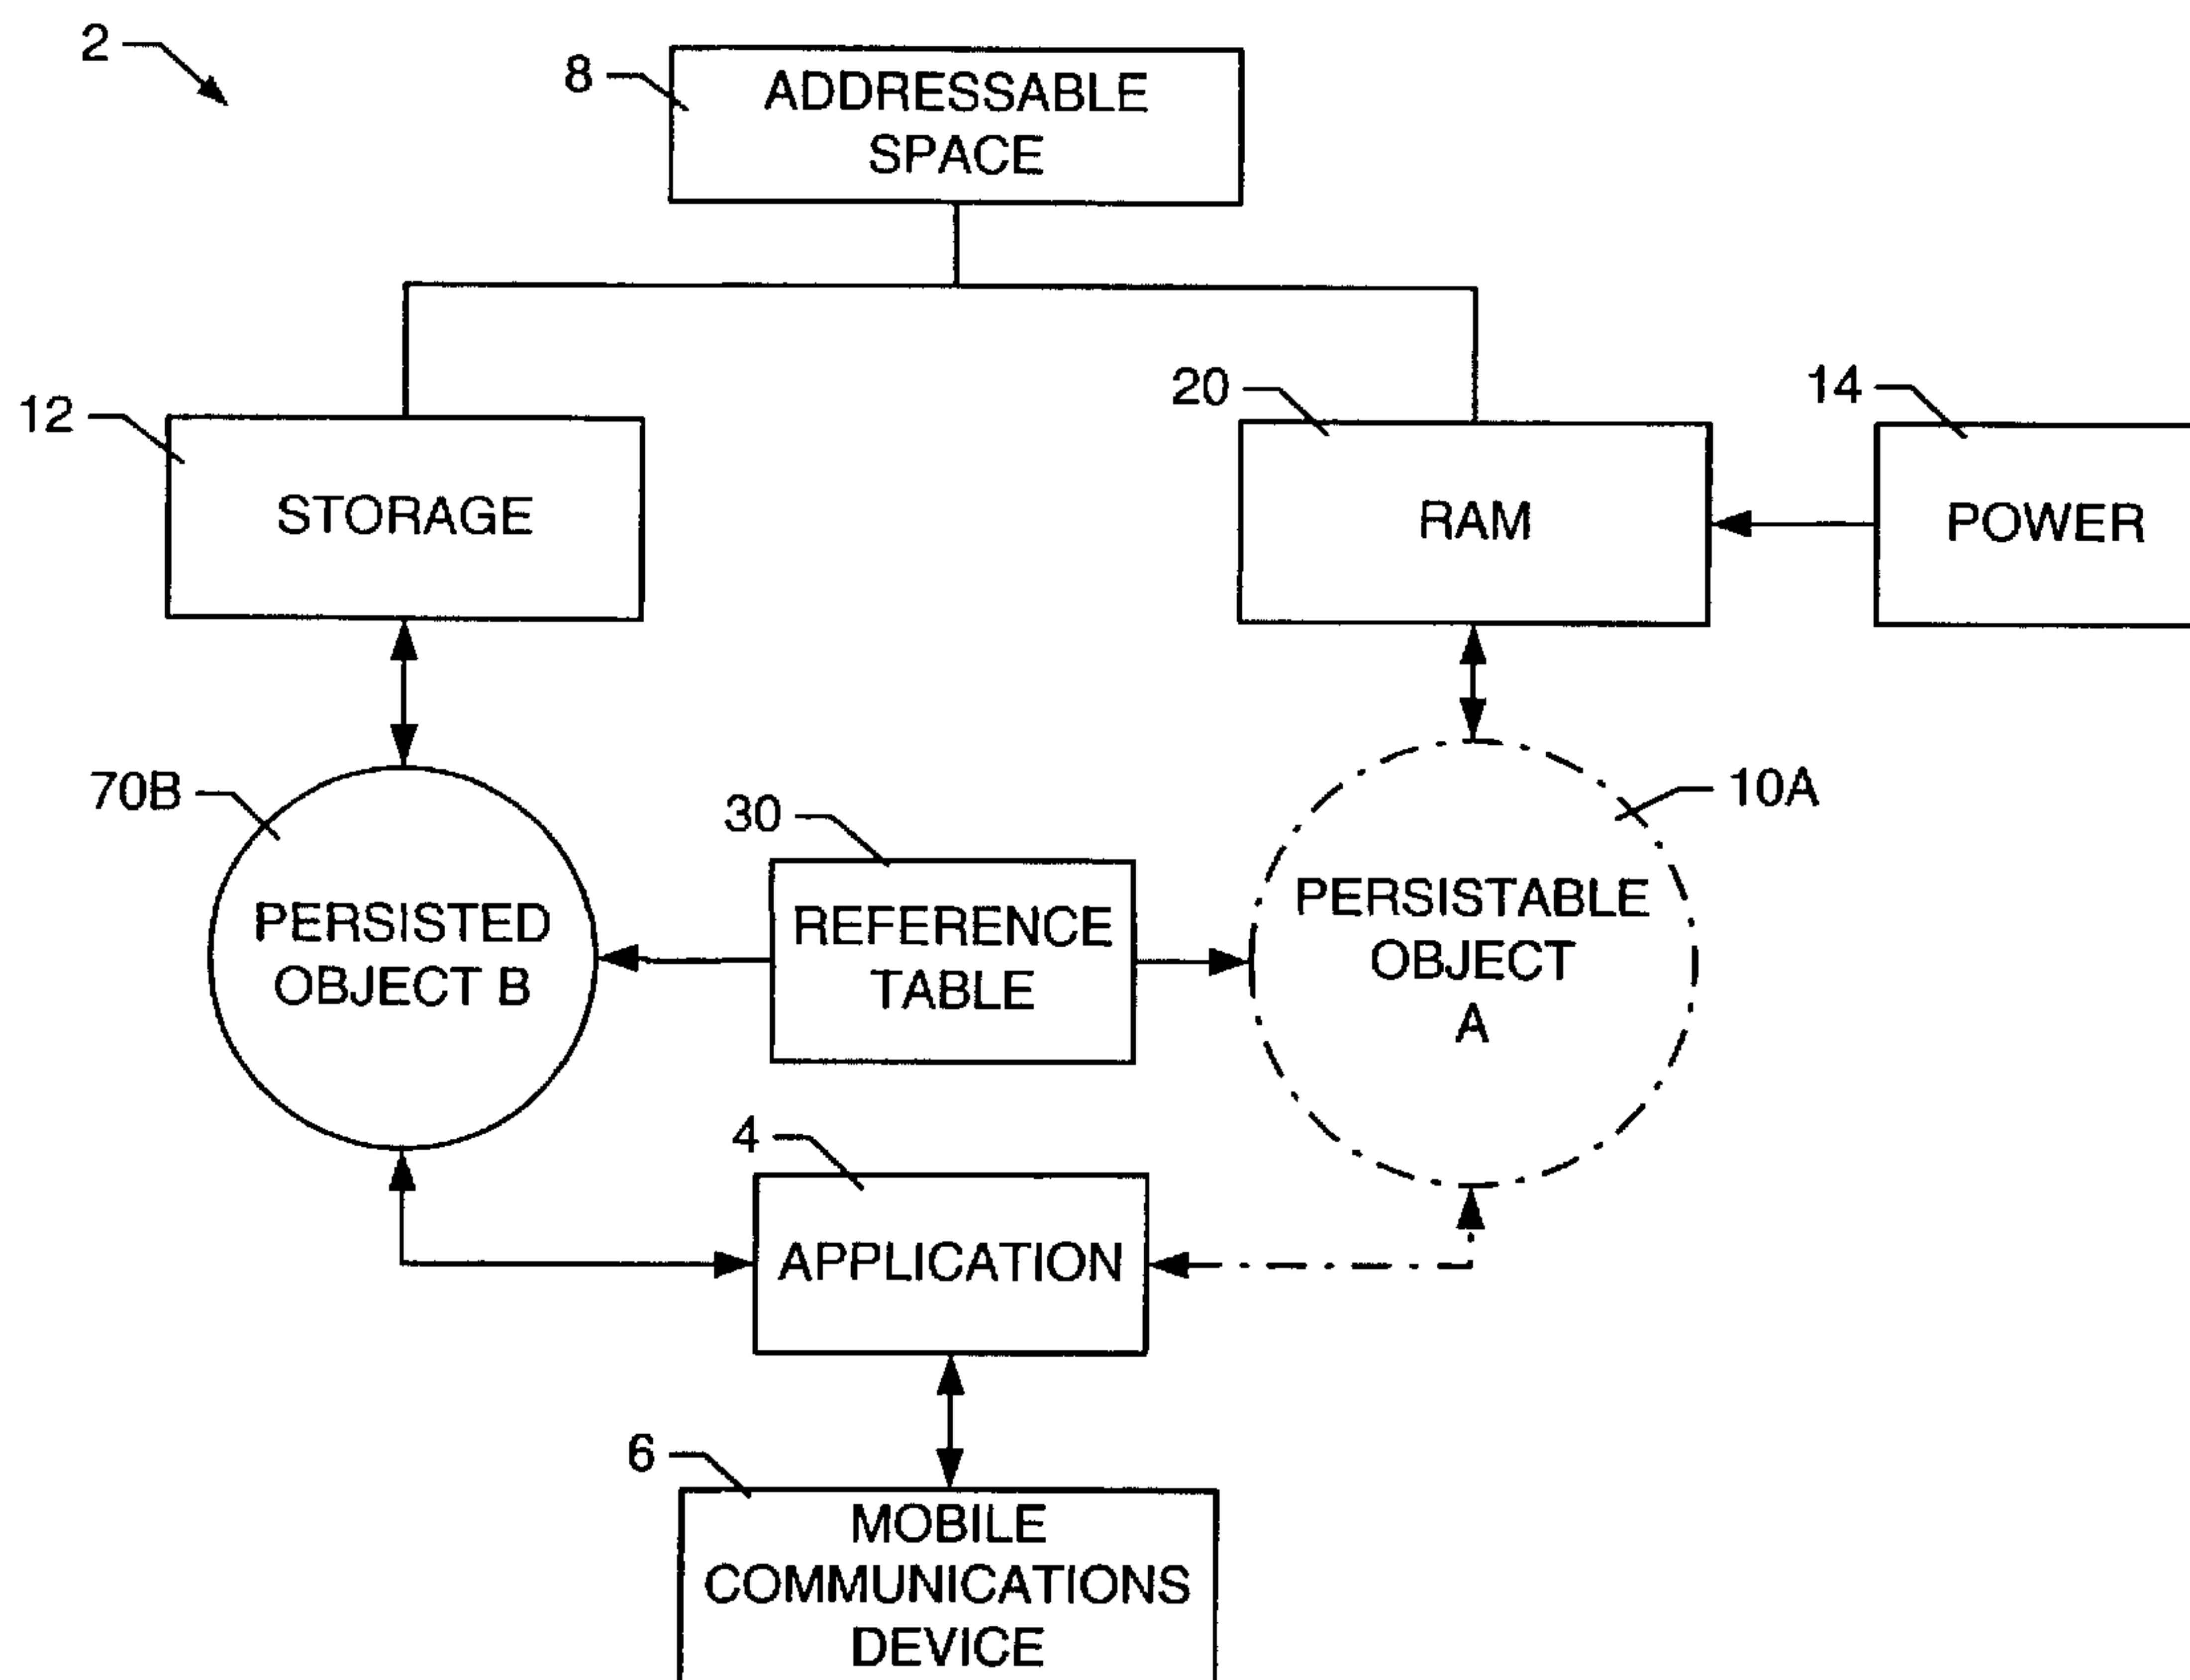




(86) Date de dépôt PCT/PCT Filing Date: 2002/07/02
 (87) Date publication PCT/PCT Publication Date: 2003/01/16
 (85) Entrée phase nationale/National Entry: 2003/12/30
 (86) N° demande PCT/PCT Application No.: CA 2002/001008
 (87) N° publication PCT/PCT Publication No.: 2003/005203
 (30) Priorité/Priority: 2001/07/03 (60/302,872) US

(51) Cl.Int.⁷/Int.Cl.⁷ G06F 9/44
 (71) Demandeur/Applicant:
RESEARCH IN MOTION LIMITED, CA
 (72) Inventeurs/Inventors:
DAHMS, JOHN F. A., CA;
YACH, DAVID P., CA
 (74) Agent: BORDEN LADNER GERVAIS LLP

(54) Titre : SYSTEME ET PROCEDE DE PERSISTANCE ORIENTEE OBJET
 (54) Title: SYSTEM AND METHOD OF OBJECT-ORIENTED PERSISTENCE



(57) **Abrégé/Abstract:**

A system and method for handling an application's access of persisted objects for use within a mobile communications device. The device includes a non-volatile storage device that has at least a portion of its storage locations addressable by the application. A reference table contains the addresses of persisted objects contained in the storage device. Via the reference table, an application can locate persisted objects stored in the storage device.

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

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
16 January 2003 (16.01.2003)

PCT

(10) International Publication Number
WO 03/005203 A3(51) International Patent Classification⁷: G06F 9/44

(21) International Application Number: PCT/CA02/01008

(22) International Filing Date: 2 July 2002 (02.07.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/302,872 3 July 2001 (03.07.2001) US

(71) Applicant (for all designated States except US): RESEARCH IN MOTION LIMITED [CA/CA]; 295 Phillip Street, Waterloo, Ontario N2L 3W8 (CA).

(72) Inventors; and

(75) Inventors/Applicants (for US only): DAHMS, John, F., A. [CA/CA]; 296 Castlefield Avenue, Waterloo, Ontario N2K 2N1 (CA). YACH, David, P. [CA/CA]; 254 Castlefield Ave, Waterloo, Ontario N2K 2N1 (CA).

(74) Agents: PATHIYAL, Krishna, K. et al.; Research In Motion Limited, 295 Phillip Street, Waterloo, Ontario N2L 3W8 (CA).

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

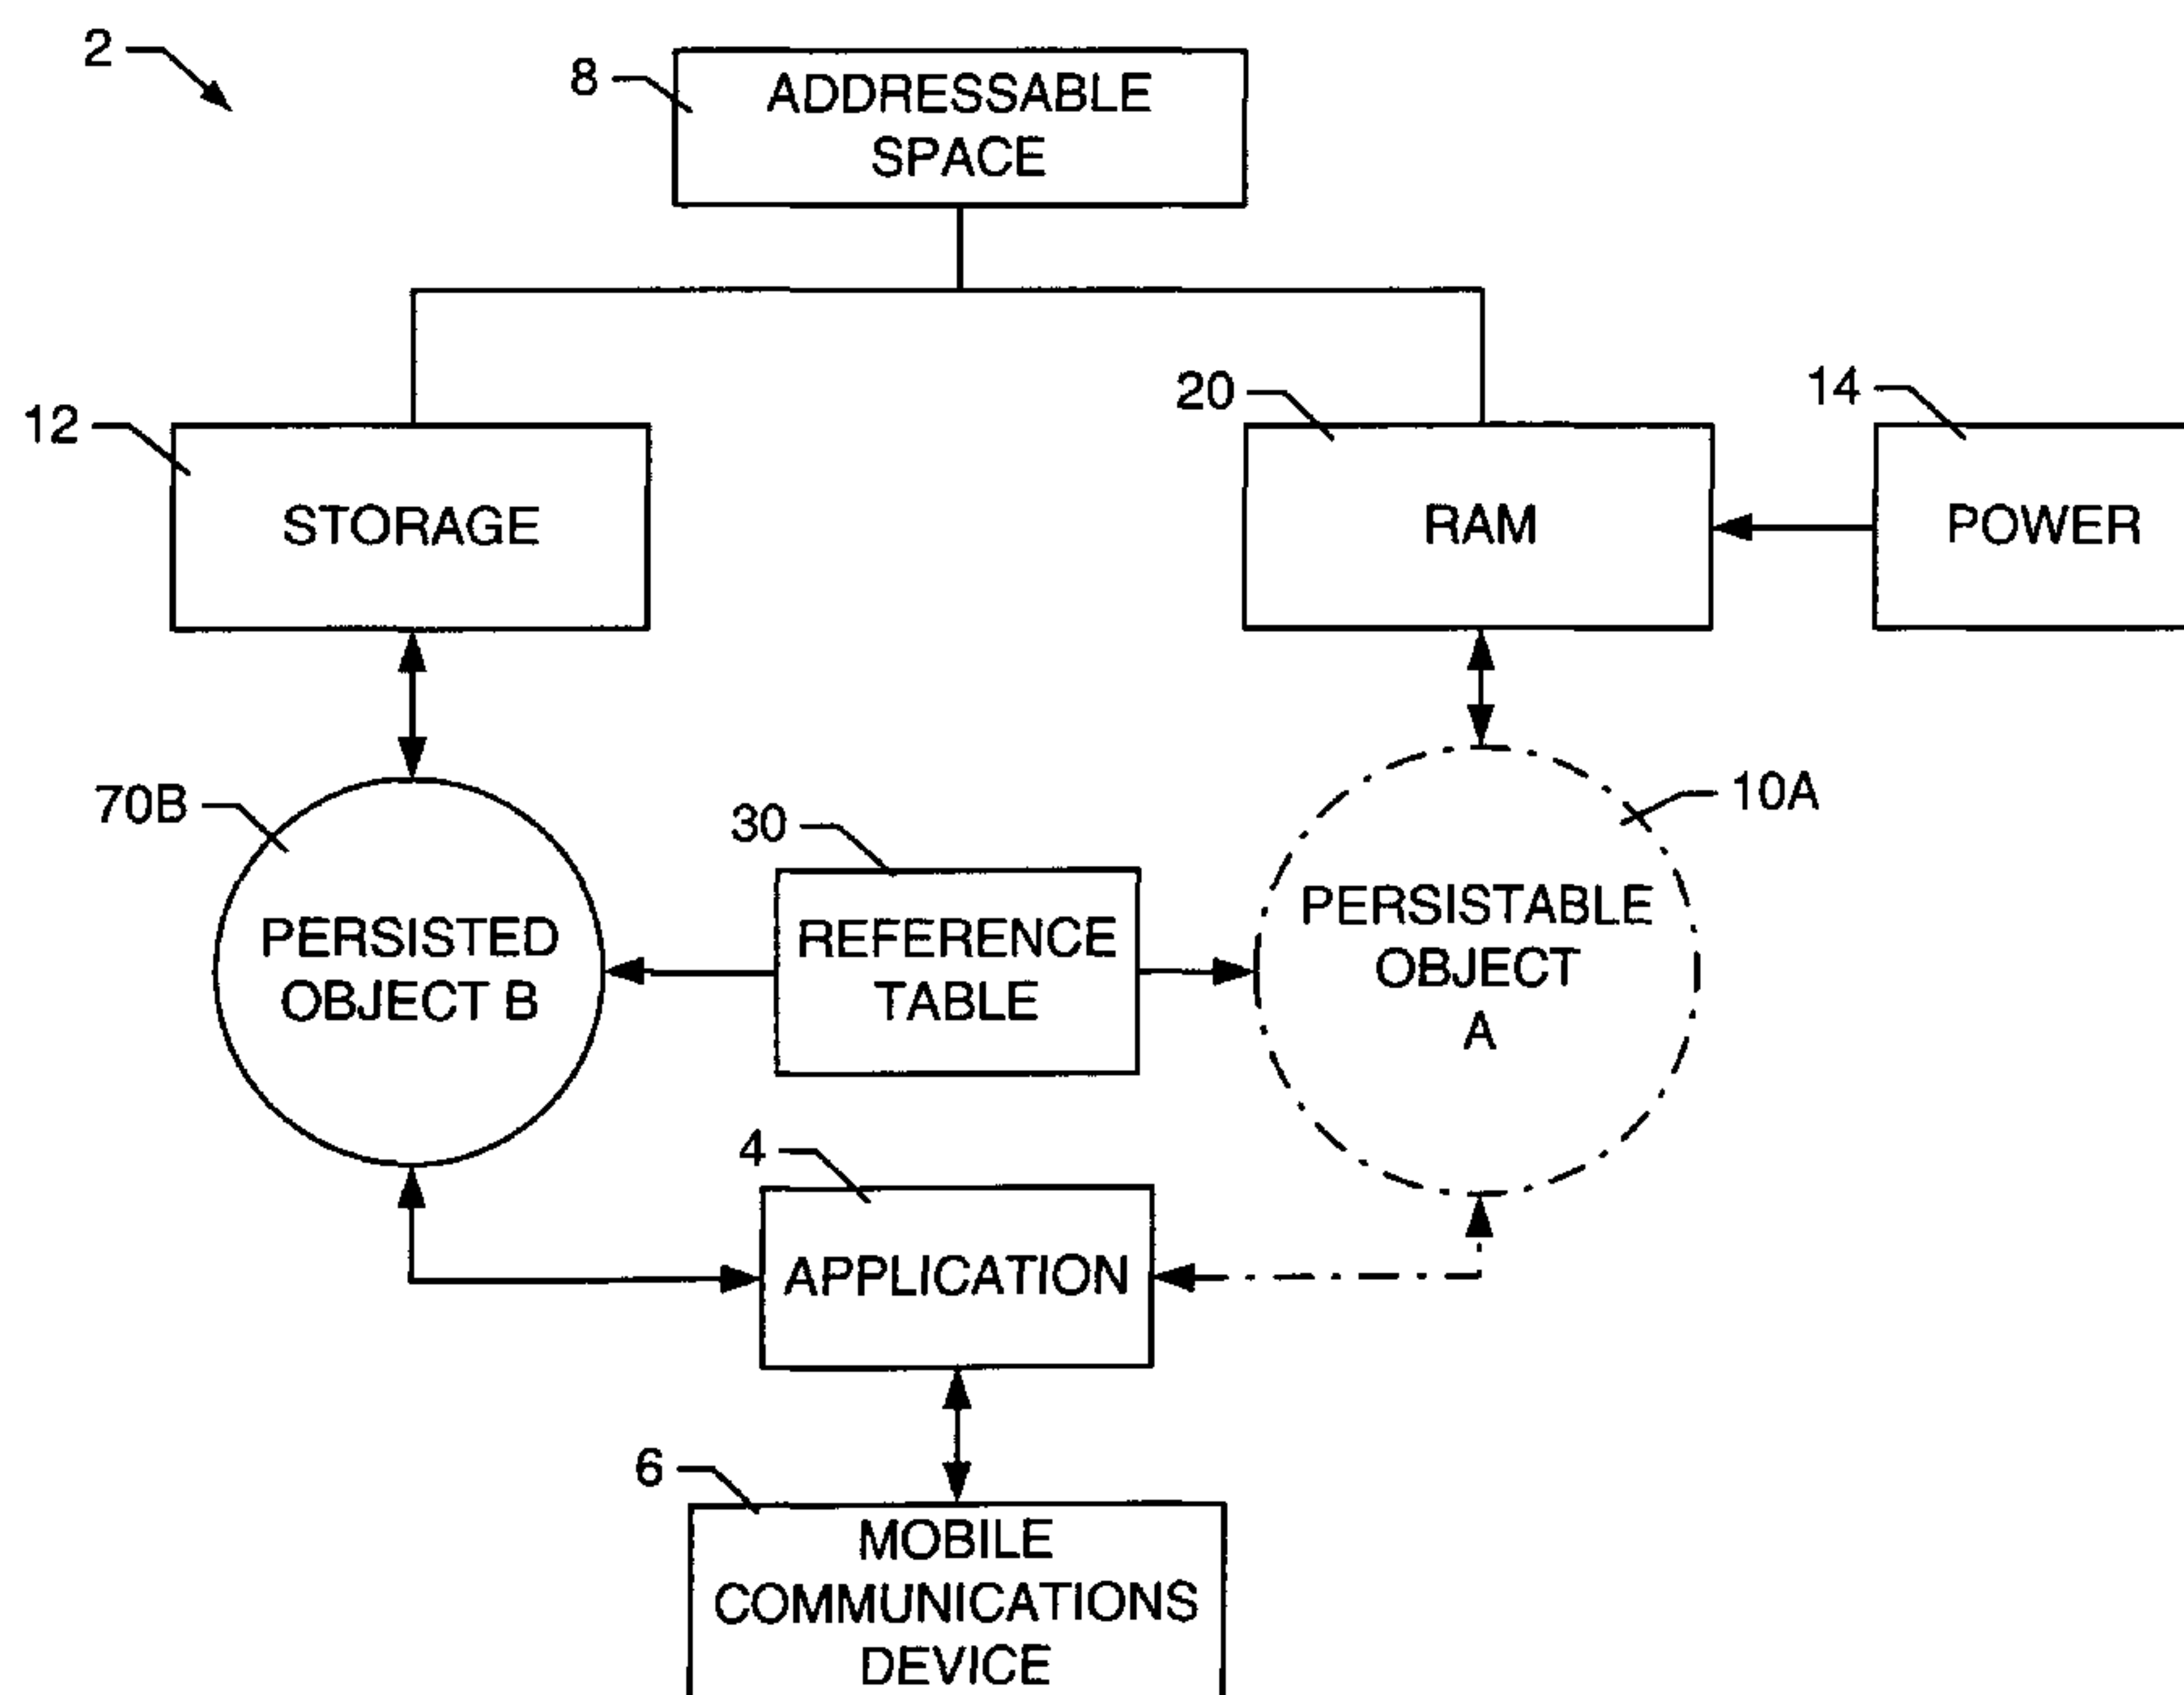
(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for the following designations AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG,

[Continued on next page]

(54) Title: SYSTEM AND METHOD OF OBJECT-ORIENTED PERSISTENCE



(57) Abstract: A system and method for handling an application's access of persisted objects for use within a mobile communications device. The device includes a non-volatile storage device that has at least a portion of its storage locations addressable by the application. A reference table contains the addresses of persisted objects contained in the storage device. Via the reference table, an application can locate persisted objects stored in the storage device.



WO 03/005203 A3



- MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW, ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii)) for the following designations AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW, ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),*
- European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)*
- *of inventorship (Rule 4.17(iv)) for US only*
- Published:**
- *with international search report*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*
- (88) Date of publication of the international search report:**
9 October 2003
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

Title

System and Method of Object-Oriented Persistence

Related Application

5 This application claims priority to U.S. provisional application
Serial No. 60/302,872 entitled "System and Method of Object-Oriented
Persistence" filed July 3, 2001. By this reference, the full disclosure, including
the drawings, of U.S. provisional application Serial No. 60/302,872 is
incorporated herein.

10 **Background**

Field of the Invention

 This invention relates to the field of information storage.
Particularly, this invention relates to storing object-oriented information within
an object-oriented environment.

15 ***Discussion of the Related Art***

 Presently known techniques for storing objects, such as the
serializable interface of the java.io package in Sun Microsystems Java™ 2 SE
(v1.3), may impose undue limitations such as assuming that the storage
involves an underlying file on a file system. For instance, known techniques
20 for storing objects may require that the object be first retrieved from storage,
then restored to a useable state, in order to be used.

Summary

 The present invention overcomes such disadvantages as well as
25 others. In accordance with the teaching of the present invention, a system
and method are disclosed that handle an application's access of persisted

objects for use within a mobile communications device. The device includes a non-volatile storage device that has at least a portion of its storage locations addressable by the application. A reference table contains the addresses of persisted objects contained in the storage device. Via the reference table, an application can locate persisted objects stored in the storage device. Further features of the invention will be described or will become apparent in the course of the following detailed description.

Brief Description of the Drawings

In order that the invention be more clearly understood, embodiments thereof will now be described in detail by way of example only, with reference to the accompanying drawings, in which:

FIG. 1 is a block diagram illustrating software and computer components used in persisting and unpersisting objects;

FIG. 2 is a data structure diagram illustrating a reference table;

FIGS. 3 and 4 are block diagrams illustrating an example of a format occupied by objects in RAM and that may be accessed via a reference table;

FIGS. 5 and 6 are block diagrams illustrating an example of a format occupied by persisted objects and that may be accessed via a reference table;

FIGS. 7 and 8 are block diagrams illustrating cooperation of objects in RAM and non-volatile storage;

FIGS. 9 and 10 are block diagrams illustrating un-grouped persistable objects;

FIG. 11 is a block diagram illustrating groupings of object references;

FIG. 12 is a data structure diagram illustrating a reference table for use in groupings of object references;

FIGS. 13 and 14 are block diagrams illustrating grouped persisted objects;

FIG. 15 is a flowchart illustrating an example of persisting objects;

FIG. 16 is a flowchart illustrating an example of object grouping;
FIG. 17 is a flowchart illustrating an example of un-persisting
objects;

FIG. 18 is a flowchart illustrating an example of object un-
5 grouping; and

FIG. 19 is a flowchart illustrating an example of re-population of
a reference table with persisted object references.

The same reference numerals are used in different figures to
refer to similar elements.

10

Detailed Description of the Preferred Embodiment

FIG. 1 depicts a system 2 to handle an application's 4 access of
objects (e.g., 10A and 70B) within the environment of a mobile
communications device 6. Through use of a reference table 30, the
15 application 4 accesses desired objects 10A and 70B without special regard to
whether the objects 10A and 70B are stored in random access memory
(RAM) 20 or a non-volatile storage device 12, respectively. RAM 20 and the
storage device 12 have addresses that are within an addressable space.

Some storage devices can be addressed like conventional RAM,
20 such as flash memory for instance. In the case where storage is addressed
like conventional RAM, a persisted object can be read directly from storage.
In this way, both reading and writing of persisted objects can be accomplished
while an object is persisted. The addressable space for RAM 20 and the
storage device 12 may comprise the same addressable space, or
25 substantially the same addressable space. In a different embodiment, the
addressable space may be different.

A persisted object (e.g., object 70B) is an object that has
achieved object-oriented persistence and is located in the storage device 12.
Similarly, an object that has the potential to be persisted (such as object 10A
30 in RAM 20) is defined as a persistable object. Persisting an object is defined
as the process of providing a persisted object from a persistable object.
Unpersisting an object is defined as the converse, the process of providing a

persistable object from a persisted object. Persisted objects are stored in the storage device 12 whose contents are unaffected if power 14 is removed (whereas the contents of volatile memory 20 are lost when power 14 is no longer supplied to it). As used in this description and in the appended claims, object-oriented persistence may include the format being occupied by an object in storage 12 being substantially similar to the format a corresponding object occupies in RAM 20. By using the same format in storage 12 as in RAM 20, only systematic operations are required to restore persisted objects to a useable state.

Since virtually no operations are required in the system 2 to restore a persisted object to a useable state, a persisted object can be used directly from the storage device. A useable persisted object may continue to be referenced by other objects residing in either RAM or storage, as well as reference other objects. As a consequence, a mobile communications device such as a PDA, cellular telephone, or data communications device can selectively move objects to and from storage without consuming significant resources. Once an object is persisted, it can continue to be used in substantially the same way as an object in RAM. This is a significant advantage over an unusable serialized object, which has to be loaded to be used, and may typically take up to three orders of magnitude more time to become useable. It is noted that file access times may be measured in milliseconds whereas memory access may be measured in nanoseconds.

For example, in a mobile device address-book application, an Address object might have a Name attribute and a Telephone attribute. Since most address-book entries do not change very often and are used regularly, it may be desirable to persist address-book entries so that they can be both stored and read in the future. Address objects benefit by being persistable, and benefit from being useable persisted objects. Since the Name and Telephone attributes (being integral to an Address object) may also be useable persisted objects, an object-oriented runtime is able to read useable persisted Address objects directly while they reside in storage 12 without need to retrieve Address objects from storage 12. This is a significant savings that may allow a persisted object to take several orders of magnitude less

time to become useable in comparison to a serialized object. Furthermore, by using substantially the same format in RAM 20 and in storage 12, persisted objects are less likely to become obsolete.

The system 2 includes a reference table 30 that locates objects 5 10A and 70B whether they are in RAM 20 or the storage device 12. For example where an application 4 wishes to access persistable object A 10A, the reference table 30 provides the memory address of the persistable object A 10A which is located in RAM 20. If the application 4 also has need to access persisted object B 70B, then the reference table 30 provides the 10 memory address of the persisted object B 70B, which is located in the storage device 12. It should be understood that the system 2 may store objects for direct access by the application 4 in both the storage device 12 and RAM 20, or solely in the storage device 12, or solely in RAM 20. The system 2 is configured with the combination that best suits the application at hand. It 15 should also be understood that one or more applications may utilize the system 2 to access the objects referenced in the reference table 30. Further it should be understood that the system 2 may include persistence of module objects to provide useable persisted modules. Modules may be used, for instance, to hold target-linked object classes. If the module is not useable 20 from storage, virtually no operations are required to render it useable once unpersisted. However, by useably persisting modules, it is possible to execute code for an object-oriented runtime directly from storage 12.

FIG. 2 depicts a data structure of the reference table 30. Indexed storage elements 32 within the reference table 30 contain addresses 25 34 to the objects stored in the storage device or in RAM. Given a desired object's index 36 to a storage element 32, the desired object's address 38 is obtained from the reference table 30. An object's address 34 in a storage element 32 reflects whether the object is in a storage device or RAM. The reference table 30 obviates an application having to specially recognize where 30 the object is stored.

FIGS. 3 and 4 depict an example of a format occupied by objects in RAM and that may be accessed via the reference table: Two objects 10A and 10B are shown as they might appear somewhere in RAM 20.

The reference table 30 is also situated in RAM 20. The reference table 30 has several storage elements 35A and 35B of a fixed size "w" 37 to simplify the indexed access to storage elements. Each used storage element 35A and 35B corresponds to an object 10A and 10B which is somewhere in an addressable space. For instance object A 10A finds correspondence with storage element index "a" 35A whereas object B 10B finds correspondence with storage element index "b" 35B. The address 40A and 40B of corresponding objects 10A and 10B is stored in a storage element 35A and 35B, so that knowing the index of an object in the reference table 30 it is possible to obtain the address 40A and 40B of an object. This is done by first obtaining the address @R 50 of the reference table 30. Then, given an object's reference, such as "a" 55A for object A 10A, the address of the storage element $@(R+a*w)$ 60A can be obtained by multiplying the index 55A "a" by the size "w" 37 of each storage element and adding to the resulting value the address @R 50 of the reference table 30. The address of the storage element $@(R+b*w)$ 60B corresponding to object B 10B can be obtained using the same technique by using index "b" in lieu of "a". Since the "a" storage element 35A holds the address of the corresponding object A 10A, resolving the contents of the storage element 35A provides the address @A 40A of object A 10A. Also shown is how each object 10A and 10B contains within its format its own reference 55A and 55B related to the reference table 30. Also shown is how, for example, object A 10A contains within its format a reference "b" 65B to object B 10B. This allows a runtime context within the scope of object A 10A to be able to access object B 10B.

Although not expressly shown in the drawings, other attributes such as references to other objects, or atomic data types, may be present within the formats of objects. Although not expressly shown in the drawings, many instances of objects may exist in an addressable space. In the example of FIG. 3, two instances 10A and 10B of classes have been illustrated. It should be understood that each instance may be of a different class, or alternatively that each instance may be of the same class. In the first case, the object A 10A may be an instance of class "A" whereas object B 10B may be an instance of class "B". For example, object A might be an instance of an

Address class and object B might be an instance of a Name class. There may be several instances of the Address class and Name class situated somewhere in the addressable space, such as in RAM. For the sake of brevity and clarity, only one instance of class "A" and "B" were discussed.

5 Alternatively in the second case, object A 10A and object B 10B could both be instances of the same class. Objects A and B are used by way of example only with the understanding that in practice there may be many objects which are instances of many classes.

FIGS. 5 and 6 depict an example of the format occupied by
 10 persisted objects and that may be accessed via a reference table. Two useable persisted objects 70A and 70B are illustrated as they might appear somewhere in storage, such as flash 80. The address of persisted objects 70A and 70B are stored in reference table 30 storage elements 35A and 35B, so that by knowing the index of a persisted object in the reference table 30 it
 15 is possible to obtain the address 90A, 90B of the persisted objects. By obtaining the address @R 50 of the reference table 30, given a persisted object's reference, such as "b" 55B for the example B persisted object 70B, the address of the storage element $@(R+b*w)$ 60B can be obtained by multiplying the index 55B "b" by the size "w" 37 of each storage element and
 20 adding the resulting value to the address @R 50 of the reference table 30. Since the "b" storage element 35B holds the address of the corresponding persisted object B 70B, resolving the contents of the storage element 35B provides the address @B 90B of persisted object B 70B. The same technique can be used to provide the address 90A of persisted object 70A from the
 25 address of storage element $@(R+a*w)$ 60A. Also shown is how each persisted object 70A and 70B contains within its format its own reference 55A and 55B related to the reference table 30. Also shown is how, for example, persisted object A 70A contains within its format a reference "b" 65B to persisted object B 70B. This allows a runtime context within the scope of
 30 persisted object A 70A to be able to access persisted object B 70B.

Persisted objects A 70A and B 70B are situated in flash 80. The format in flash 80 of persisted objects 70A, 70B of FIG. 5 is substantially similar to corresponding objects 10A and 10B in RAM 20 of FIG. 3. Since

flash 80 is addressable and can be read substantially in the same way as RAM 20, and since the persisted objects 70A, 70B are substantially in the same format as corresponding un-persisted objects 10A and 10B of FIG. 3, persisted objects 70A, 70B can be referenced by the runtime context "as if" they were in RAM 20, thereby making them useable persisted objects.

When power is shut off, persisted objects A 70A and B 70B will remain stored in flash 80. However, reference table 30 will be lost making persisted objects A 70A and B 70B temporarily unusable. When power is restored, reference table 30 can be systematically re-populated by reading flash 80 since every persisted object 70A, 70B has within its format its own reference, thereby restoring useable persisted objects A 70A and B 70B. Re-population will be described in greater detail in reference to FIG. 19.

In reference to FIGS. 7 and 8, a block diagram illustrating the cooperation of objects in RAM and persisted objects is described. An object 10A and a useable persisted object B 70B are illustrated, as they might appear somewhere in RAM 20 and flash 80 respectively. Also illustrated is a reference table 30, situated somewhere in RAM 20. This allows a runtime context within the scope of object A 10A to be able to access usable persisted object B 70B, regardless of the fact that object B 70B is in flash 80 and object A 70A is in RAM 20.

A persisted object, such as object B 70B, may contain a reference (not shown) to a non-persisted object, such as object A 10A. This allows a runtime context within the scope of persisted object B 70B to be able to access object A 10A, regardless of the fact that object B 70B is in flash 80 and object A 70A is in RAM 20. Furthermore, a runtime context within the scope of persisted object B 70B, such as an instance method context, may receive an object reference to object A 10A as a parameter.

In reference to FIGS. 9 and 10, block diagrams illustrating ungrouped persistable objects are described. Two objects 10A and 10B are illustrated as they might appear somewhere in RAM 20. Also illustrated is a reference table 30, situated somewhere in RAM 20. In situations where it is desired to impose an upper limit on the amount of RAM 20 dedicated to storing object references, reference table 30 may have a maximum of n

storage elements. This might be the case, for instance, in a device with limited RAM 20. By imposing an upper limit n , the reference table 30 can be maintained within a reasonable size. Each used storage element corresponds to objects 10A and 10B which are somewhere in an addressable space, which makes un-grouped objects 10A and 10B compatible with other objects thus far described. For instance object A 10A finds correspondence with storage element "a" 35A whereas object B 10B finds correspondence with storage element "b" 35B. The address of corresponding objects 10A and 10B are stored in storage element 35A, 35B, so that knowing the index of an object in the reference table 30 it is possible to obtain the address 40A, 40B of an object. This is done by first obtaining the address @R 50 of the reference table 30. Then, given an objects reference 55A, such as "0 | a" for the example A object 10A, the address of the storage element $@(R+a*w)$ 60A can be obtained by multiplying the index 55A "a" by the size "w" 37 of each storage element and adding to the result the address @R 50 of the reference table 30. Similarly, the address of storage element $@(R+b*w)$ 60B can be obtained given an objects reference 55B, such as "0|b" for the example B object B 10B. The address of the object @A 40A is contained in the corresponding storage element 35A. Similarly, the address of the object @B 40B is contained in the corresponding storage element 35B. The purpose of the "offset" is described further in FIG. 11 in reference to grouped persisted objects. Since both objects A 10A and B 10B are un-grouped, the "offset" value is 0 for each, as is illustrated by the corresponding "this ref= 0|a" for object A 10A and "this ref=0|b" for object B 10B. Also shown is how flash 80 might comprise several blocks 120, of which only 2 are illustrated for brevity. This is meant to illustrate how some flash 80 may be addressed to be read in a manner substantially similar to RAM 20, but may be only written to in blocks 120.

The block technique may also use object grouping. Grouping conserves reference table's RAM 20 usage. FIG. 11 shows groupings 102 of object references so that objects integral to a "root" object are located in addressable space in substantially close proximity to the "root" object. In this way, RAM 20 can be used more efficiently since it only uses a single object

reference element 104 in the reference table 30 for each "root" object in a group. Grouped object references 102 are updated so that they are offset relative to the "root" object. Object reference grouping is defined as relating object addresses in a group of objects by an offset to a "root" object reference
 5 that integrates the group.

FIG. 12 depicts a reference table data structure for use in groupings of object references. Indexed storage elements 32 within the reference table 30 contain addresses 104 to the "root" objects stored in the storage device or in RAM. Given a desired object's root object's index 106,
 10 the root object's address is obtained from the reference table 30. A root object's address in a storage element 32 reflects whether the object is in a storage device or RAM. The "root" object 108 contains object references to its integral grouped objects, such as the desired object. The desired object is referenced in the format of the "root" object. The address 110 of the desired
 15 object is related to the address of the "root" object by an "offset" as discussed in greater detail in reference to FIGS. 13 and 14.

In reference to FIGS. 13 and 14, block diagrams illustrating grouped persisted objects are described. Two persisted objects 70A, 70B are illustrated as they might appear in block 120 of flash 80. Also illustrated is a
 20 reference table 30, situated somewhere in RAM 20. The reference table 30 has several storage elements 35 of a fixed size "w" 37 for the indexed access to storage elements. Each used storage element, such as "a" 35A, corresponds to an object 70A which is somewhere in an addressable space. For instance object A 70A finds correspondence with storage element index
 25 "a" 35A. However, object B 70B does not find immediate correspondence with storage element index "b" – this is because object B 70B has been grouped with object A 70A in order to conserve reference table 30 RAM 20 usage. Note how object A 70A is still referenced in the reference table 30, and that its reference 55A "0 | a" has an "offset" of 0. Object A 70A is therefore a "root"
 30 object, harboring object references to its integral grouped objects, such as object B 70B. Object B 70B is referenced in the format of the "root" object 70A, the address @B 90B of object B 70B is related to the address @A 90A of "root" object A 70A by an "offset". The "offset" 130 is the difference between

the flash address @A 90A of object A and the flash address @B 90B of object B. Note that all objects in the group 70A, 70B are in the same block 120 of flash 80. The range of offset 130 is related to the size of block 120. For instance, if flash 80 uses a block 120 of size 64K, then the offset can be represented using 16 bits. The address of "root" object 70A is stored in storage element 35A so that by knowing the index of a "root" object 70A in the reference table 30 it is possible to obtain the address @A 90A of a "root" object. Resolving the address @B 90B of object B 70B can still be related to the reference table 30 via "root" object A 70A. By adding to the address of reference table 30, @R 50 index "a" 35A (multiplied by the width 37) provides the address of the storage element $(R+a*w)$ 60A (which contains the address @A 90A of the "root" object A 70 A). From the "root" address @A 90A, by adding the "offset" 130 of object B's 70B reference "offset | a" 65B in the format of object A 70A, the resulting address $(A+ \text{offset})$ 60B provides address @B 90B of the grouped object B 90B. Also shown is how each object 70A, 70B contains within its format its own reference 55A, 55B ultimately still accessible via reference table 30 for "root" objects, and via "root" objects for grouped objects.

With respect to the size of block 120 of flash 80, the number of reference table Elements "n", and the size of each reference table 30 storage element "w" 37 can be adapted to suit particular needs. For example, if w is 32 bits, then a 32-bit addressable space may be used. If a flash block size of 64K is used, then 16 bits of object references 55A, 55B 65B may be used for the offset and the remaining bits may be used as an index into the reference table 30, as illustrated with the "offset | i" notation. The maximum number of elements "n" depends on the number of object references desired, and therefore imposes the range of the lower order bits representing the index "i" in object references. Of interest is using 16 bits for the index thereby providing a 32-bit object reference, comparable in size to the addressable space address size.

An example of potential RAM 20 savings is given next assuming the values in the above paragraph. The example will illustrate the positive combination of several of the techniques taught thus far. Consider an

embodiment within a mobile e-mail device. Assume that an Address object instance might have 9 object references, whereas an Email object instance might have 24 object references. To reference 5000 Address objects and 2000 Email objects simultaneously, without using grouped objects, then
5 $100,000=(5000 \times 10+2000 \times 25)$ object references are needed. A minimum of $\sim 400\text{K}$ of RAM would be required for a reference table. However, consider the situation where each of the 5000 Address objects and each of the 2000 Email objects are "root" objects, and each of the 9 object references for each Address object and each of the 24 object references for each Email object are
10 grouped and "offset" from each group's "root" Address or Email object. In this grouped form only $7000=(5000+2000)$ object references for the "root" objects need to be stored in the reference table. A minimum of $\sim 28\text{K}$ of RAM would be required for a reference table – this represents a significant 93% saving as compared to the un-grouped case. If only 400K of RAM is actually available,
15 non-grouped 5000 Address objects and 2000 Email objects would quickly cause a shortage of object references, whereas the scenario using grouped objects leaves ample room to continue allocating new object references. Furthermore, grouped objects may further conserve RAM 20 usage by being persisted to storage, such as flash 80. Further still, by making the "root"
20 persisted objects usable directly from flash 80, Addresses and Emails can be read and used directly from flash without consuming substantial RAM 20. If flash 80 is block-writable, then when Addresses and Emails are grouped and persisted, they can be considered "read-only". In order to modify a "read-only" object, the object is un-persisted and optionally un-grouped, modified,
25 and optionally grouped and persisted anew.

In situations where flash cannot be written to as easily as it is read, then object grouping provides the additional indication that grouped persistent object can be considered "read-only".

It is within the scope of this invention that "root" objects reside in
30 RAM. The use of a "block" illustrates that the addresses within a "block" are substantially proximate, regardless of whether the "block" is in RAM 20 or storage 80. Removing the limitation of "block" writing-only flash 80 still allows for the use of object grouping and un-grouping, and therefore the use of

grouping for all manner of objects is within the scope of this invention, whether a grouped object is persisted, or is in RAM. An object need not be persistable to be groupable, nor need be groupable to be persistable.

In reference to FIG. 15, a flowchart illustrating persisting objects is described. At step 210, a persistable object's 10 format is read from memory 20. At step 220, a substantially similar persisted object 70 format is written to storage 80:

Optionally, at step 230 the reference table 30 is updated to reflect the change of address from memory address @Ar 40 to storage address @Af 90. This makes the persisted object 70 useable in substantially the same way as the persistable object 10, thereby allowing the persistable object 10 to be destroyed and the persisted object 70 to transparently take its place.

Optionally, at step 300 the objects referenced in the format of the persisted object 70 are persisted in substantially close proximity to the persisted object, as is illustrated in further detail next in reference to FIG. 16.

In reference to FIG. 16, a flowchart illustrating object grouping 300 is described. At step 200, a persistable object 10B referenced in a root object 70A is persisted, for instance according to the steps described in FIG. 15. The persisted object 70B is located substantially proximate to the root object 70A, and offset 130 from it so as to facilitate grouping.

At step 310, the root object 70A reference is updated to point to the persisted object of step 200, by using a reference which is offset from the root object "offset | a". Optionally, the reference of the persisted object 70B within its format is also updated to "offset | a". Optionally, at step 320 the reference in reference table 30 to the grouped object 70B is removed. The grouped object 70B is still useable via root object 70A.

In reference to FIG. 17, a flowchart illustrating un-persisting objects 400 is described. At step 410, a persisted object's 70 format is read from storage. At step 420, a substantially similar persistable object 10 format is written to memory 20.

Optionally, at step 430 the reference table 30 is updated to reflect the change of address from storage address @Af 90 to memory

address @Ar 40. This makes the persistable object 10 useable in substantially the same way as the persisted object 70, thereby allowing the persisted object 70 to be destroyed and the persistable object 70 to transparently take its place.

5 Optionally, at step 500 the objects referenced by offset, i.e. grouped in the format of the persistable object 10 are ungrouped, as is illustrated in further detail next in reference to FIG. 18.

 In reference to FIG. 18, a flowchart illustrating object ungrouping 500 will be described presently. At step 400, a grouped object 70B
10 referenced from a root object 70A is unpersisted, for instance according to the steps of FIG. 17. After step 400, the grouped object 70B continues to exist contemporaneously with the ungrouped 10B counterpart, as is illustrated by the presence of the @Af 90A reference in reference table 30. Ungrouping
15 may create new copies of all objects in the group, relocate all the references between those objects, and return a new root object. Since reference to grouped objects may still be in use when an object is ungrouped, grouped
 references continue to resolve properly to grouped objects for as long as the references exist.

 Optionally, at step 510, when it is ascertained that grouped
20 references are no longer in use, the grouped objects can be garbage collected thereby recovering the resources used. This is illustrated by the absence of reference @Af 90A in reference table 30. When garbage
 collection has taken place, only the grouped object 70B is destroyed whereas ungrouped object 10B may still exist. Garbage collection for root object 90A
25 and grouped object 90B can be triggered, for instance by explicitly setting all grouped object references to null, and setting all root object references to null.

 In reference to FIG. 19, a flowchart illustrating the re-population of a reference table with persisted object references in one embodiment is described. At step 610, a reference table 30 is created in memory 20. At step
30 620, at least one persisted object 70 reference is added to the reference table 30. Since each persisted object's format contains its own reference, it is possible to retrieve the reference directly from storage without having an entry in the reference table.

It will be appreciated that the above description relates to the invention by way of example only. Many variations on the invention will be obvious to those knowledgeable in the field, and such obvious variations are within the scope of the invention as described and claimed, whether or not
5 expressly described.

It is claimed:

1. A system for handling an application's access of objects for use within a mobile communications device, comprising:

5 random access memory (RAM) having storage locations addressable within an addressable space, said RAM containing a persistable first object having an address within the addressable space;

a non-volatile storage device having storage locations addressable within an addressable space, said storage device containing a
10 persisted second object having an address within the storage device's addressable space;

a reference table that contains the address of the persistable first object contained in the RAM and the address of the persisted second object contained in the storage device,

15 wherein the addresses of the persistable first object and the persisted second object are obtained from the reference table so that the persistable first object and persisted second object may be used by the application.

20 2. The system of claim 1 wherein the persisted second object is unpersisted into RAM, wherein format of the persisted second object and format of the unpersisted second object are substantially similar.

25 3. The system of claim 1 wherein the persisted second object is in a usable state by the application directly from the storage device.

30 4. The system of claim 3 wherein access time associated with access of the persisted second object by the application via the reference table is at least an order of magnitude less than access of an object having been persisted through serialization.

5. The system of claim 1 wherein reading and writing of the persisted second object occurs while the second persisted object is persisted.
6. The system of claim 1 wherein the address space associated with the
5 RAM overlaps at least a portion of the address space associated with the storage device.
7. The system of claim 1 wherein the address of the persisted second object is obtained via the reference table based upon an index associated with the
10 persisted second object.
8. The system of claim 7 wherein the address of the persistable first object is obtained via the reference table based upon an index associated with the persistable first object.
15
9. The system of claim 8 wherein the reference table contains storage elements having a width,
wherein the addresses of the persistable first object and persisted second object are obtained via the reference table based upon the
20 objects' respective indexes and upon the width of the storage elements.
10. The system of claim 1 wherein the persisted second object contains within its object format a reference to itself.
- 25 11. The system of claim 10 wherein the persisted second object contains within its object format a reference to a persisted third object.
12. The system of claim 11 wherein the persisted third object contains within its object format a reference to the persisted second object.
30
13. The system of claim 10 wherein the persisted second object contains within its object format a reference to the persistable first object.

14. The system of claim 13 wherein the persistable first object contains within its object format a reference to the persisted second object.
15. The system of claim 10 wherein the reference contained within the
5 persisted second object is used to repopulate address data within the reference table.
16. The system of claim 1 wherein module objects are persisted in the storage device and whose addresses are stored in the reference table,
10 wherein the addresses of the persisted module objects are obtained from the reference table so that the persisted module objects may be used by the application.
17. The system of claim 16 wherein the module objects hold target-linked
15 object classes.
18. The system of claim 17 wherein code for an object-oriented runtime is executable directly from the storage device via the reference table.
- 20 19. The system of claim 1 wherein the application comprises an e-mail address-book application.
20. The system of claim 1 wherein the mobile communications device is a personal digital assistant.
25
21. The system of claim 1 wherein the mobile communications device is a cellular telephone.
22. The system of claim 1 wherein the mobile communications device is a
30 data communications device.
23. The system of claim 1 wherein the storage device is flash memory.

24. The system of claim 23 wherein writing to the flash memory is performed in blocks, wherein a persisted object is unpersisted prior to modification, and then the unpersisted object is persisted back to the flash memory.
- 5 25. The system of claim 1 wherein data stored in the storage device remains when electrical power is removed from the storage device.
26. The system of claim 1 wherein the reference table contains a root object reference to a root object contained in the storage device.
- 10 27. The system of claim 26 wherein at least one object is stored in the storage device whose address is determined based upon an offset and the address of the root object.
- 15 28. The system of claim 26 wherein a group of objects is stored in the storage device whose addresses are determined based upon an offset and the address of the root object.
- 20 29. The system of claim 28 wherein the reference table contains a reference to objects in the group only with a reference to the root object of the group.
30. The system of claim 28 wherein the storage device comprises divisions of blocks, wherein range of the offset is based upon size of the blocks.
- 25 31. The system of claim 28 wherein the storage device is block-writable, wherein modification to the persisted second object is performed via a persistable second object that had been generated based upon the persisted second object.
- 30 32. The system of claim 31 wherein the persistable second object is ungrouped prior to the modification.
33. The system of claim 28 wherein the root object resides in RAM.

34. The system of claim 28 further comprising:

grouping means for grouping objects contained in RAM.

5 35. The system of claim 28 wherein the grouping means groups objects contained in RAM.

36. A method for handling an application's access of objects for use within a mobile communications device, comprising the steps of:

10 reading a first persistable object from random access memory (RAM);

writing the first persistable object into a non-volatile storage device at an address so as to create a first persisted object;

storing the address in a reference table; and

15 obtaining the address from the reference table so that the first persisted object may be retrieved from the storage device and used by the application, said first persisted object being in a usable state by the application directly from the storage device.

20 37. The method of claim 36 further comprising the steps of:

unpersisting the first persisted object into RAM at a second address so as to create the first persistable object; and

25 updating the reference table to reflect that the first persistable object is located at the second address, wherein format of the first persisted object and format of the first unpersisted object are substantially similar.

38. The method of claim 37 further comprising the step of:

30 obtaining the second address from the reference table so that the first persistable object may be retrieved from the RAM and used by the application.

39. The method of claim 38 further comprising the step of:

obtaining a third address from the reference table so that a second persisted object may be retrieved from the storage device and used by the application, said second persisted object being in a usable state by the application directly from the storage device.

5

40. The method of claim 38 wherein access time associated with access of the first persisted object by the application via the reference table is at least an order of magnitude less than access of an object persisted through serialization means.

10

41. The method of claim 36 wherein address space associated with the RAM overlaps at least a portion of address space associated with the storage device.

15

42. The method of claim 36 wherein the address of the first persisted object is obtained via the reference table based upon an index associated with the persisted object.

20

43. The method of claim 36 wherein the reference table contains address storage elements having a width,
wherein the address of the first persisted object is obtained via the reference table based upon the first persisted objects' index and upon the width of at least one of the storage elements.

25

44. The method of claim 36 wherein the first persisted object contains within its object format a reference to itself.

30

45. The method of claim 44 wherein the first persisted object contains within its object format a reference to a second persisted object.

46. The method of claim 44 wherein the reference contained within the first persisted object is used to repopulate address data within the reference table.

47. The method of claim 36 further comprising the step of:
persisting module objects in the storage device, said module
objects having addresses stored in the reference table, wherein the
addresses of the persisted module objects are obtained from the reference
5 table so that the persisted module objects may be used by the application.
48. The method of claim 47 wherein the module objects hold target-linked
object classes.
- 10 49. The method of claim 36 wherein the mobile communications device is a
personal digital assistant.
50. The method of claim 36 wherein the mobile communications device is a
cellular telephone.
- 15 51. The method of claim 36 wherein the mobile communications device is a
data communications device.
52. The method of claim 36 wherein the storage device is flash memory.
- 20 53. The method of claim 52 wherein writing to the flash memory is performed
in blocks, wherein a persisted object is unpersisted prior to modification, and
then the unpersisted object is persisted back to the flash memory.
- 25 54. The method of claim 36 wherein data stored in the storage device
remains when electrical power is removed from the storage device.
55. The method of claim 36 wherein the reference table contains a root object
reference to a root object contained in the storage device.
- 30 56. The method of claim 55 wherein at least one object is stored in the
storage device whose address is determined based upon an offset and the
address of the root object.

57. The method of claim 55 wherein a group of objects is persisted in the storage device and whose addresses are determined based upon an offset and the address of the root object.

5

58. The method of claim 57 wherein the reference table contains a reference to objects in a group only with a reference to the root object of the group.

59. The method of claim 57 wherein the storage device comprises divisions of blocks, wherein range of the offset is based upon size of the blocks.

10

60. The method of claim 57 wherein the root object resides in RAM.

61. The method of claim 57 further comprising the steps of:

15

unpersisting the group of objects; and
performing garbage collection upon at least one of the
unpersisted object.

62. A system for handling an application's access of objects for use within a mobile communications device, comprising:

20

a non-volatile storage device having storage locations addressable within an addressable space, said storage device containing a persisted object having an address within the addressable space;

25

a reference table that contains the address of the persisted object contained in the storage device,

wherein the address of the persisted object is obtained from the reference table so that the persisted object may be used by the application, said persisted object being in a usable state by the application directly from the storage device.

30

63. A system for handling an application's access of objects, comprising:

storage means having storage locations addressable within an addressable space, said storage means containing an object having an address within the addressable space;

a reference table that contains the address of the object

5 contained in the storage means,

wherein the address of the object is obtained from the reference table so that the object may be used by the application, said object being in a usable state by the application directly from the storage means.

10 64. The system of claim 63 wherein the storage means comprises non-volatile storage means.

65. The system of claim 64 wherein the storage means comprises volatile storage means.

15

66. The system of claim 63 wherein the storage means comprises volatile storage means.




FIG. 1




FIG. 2




FIG. 3




FIG. 4

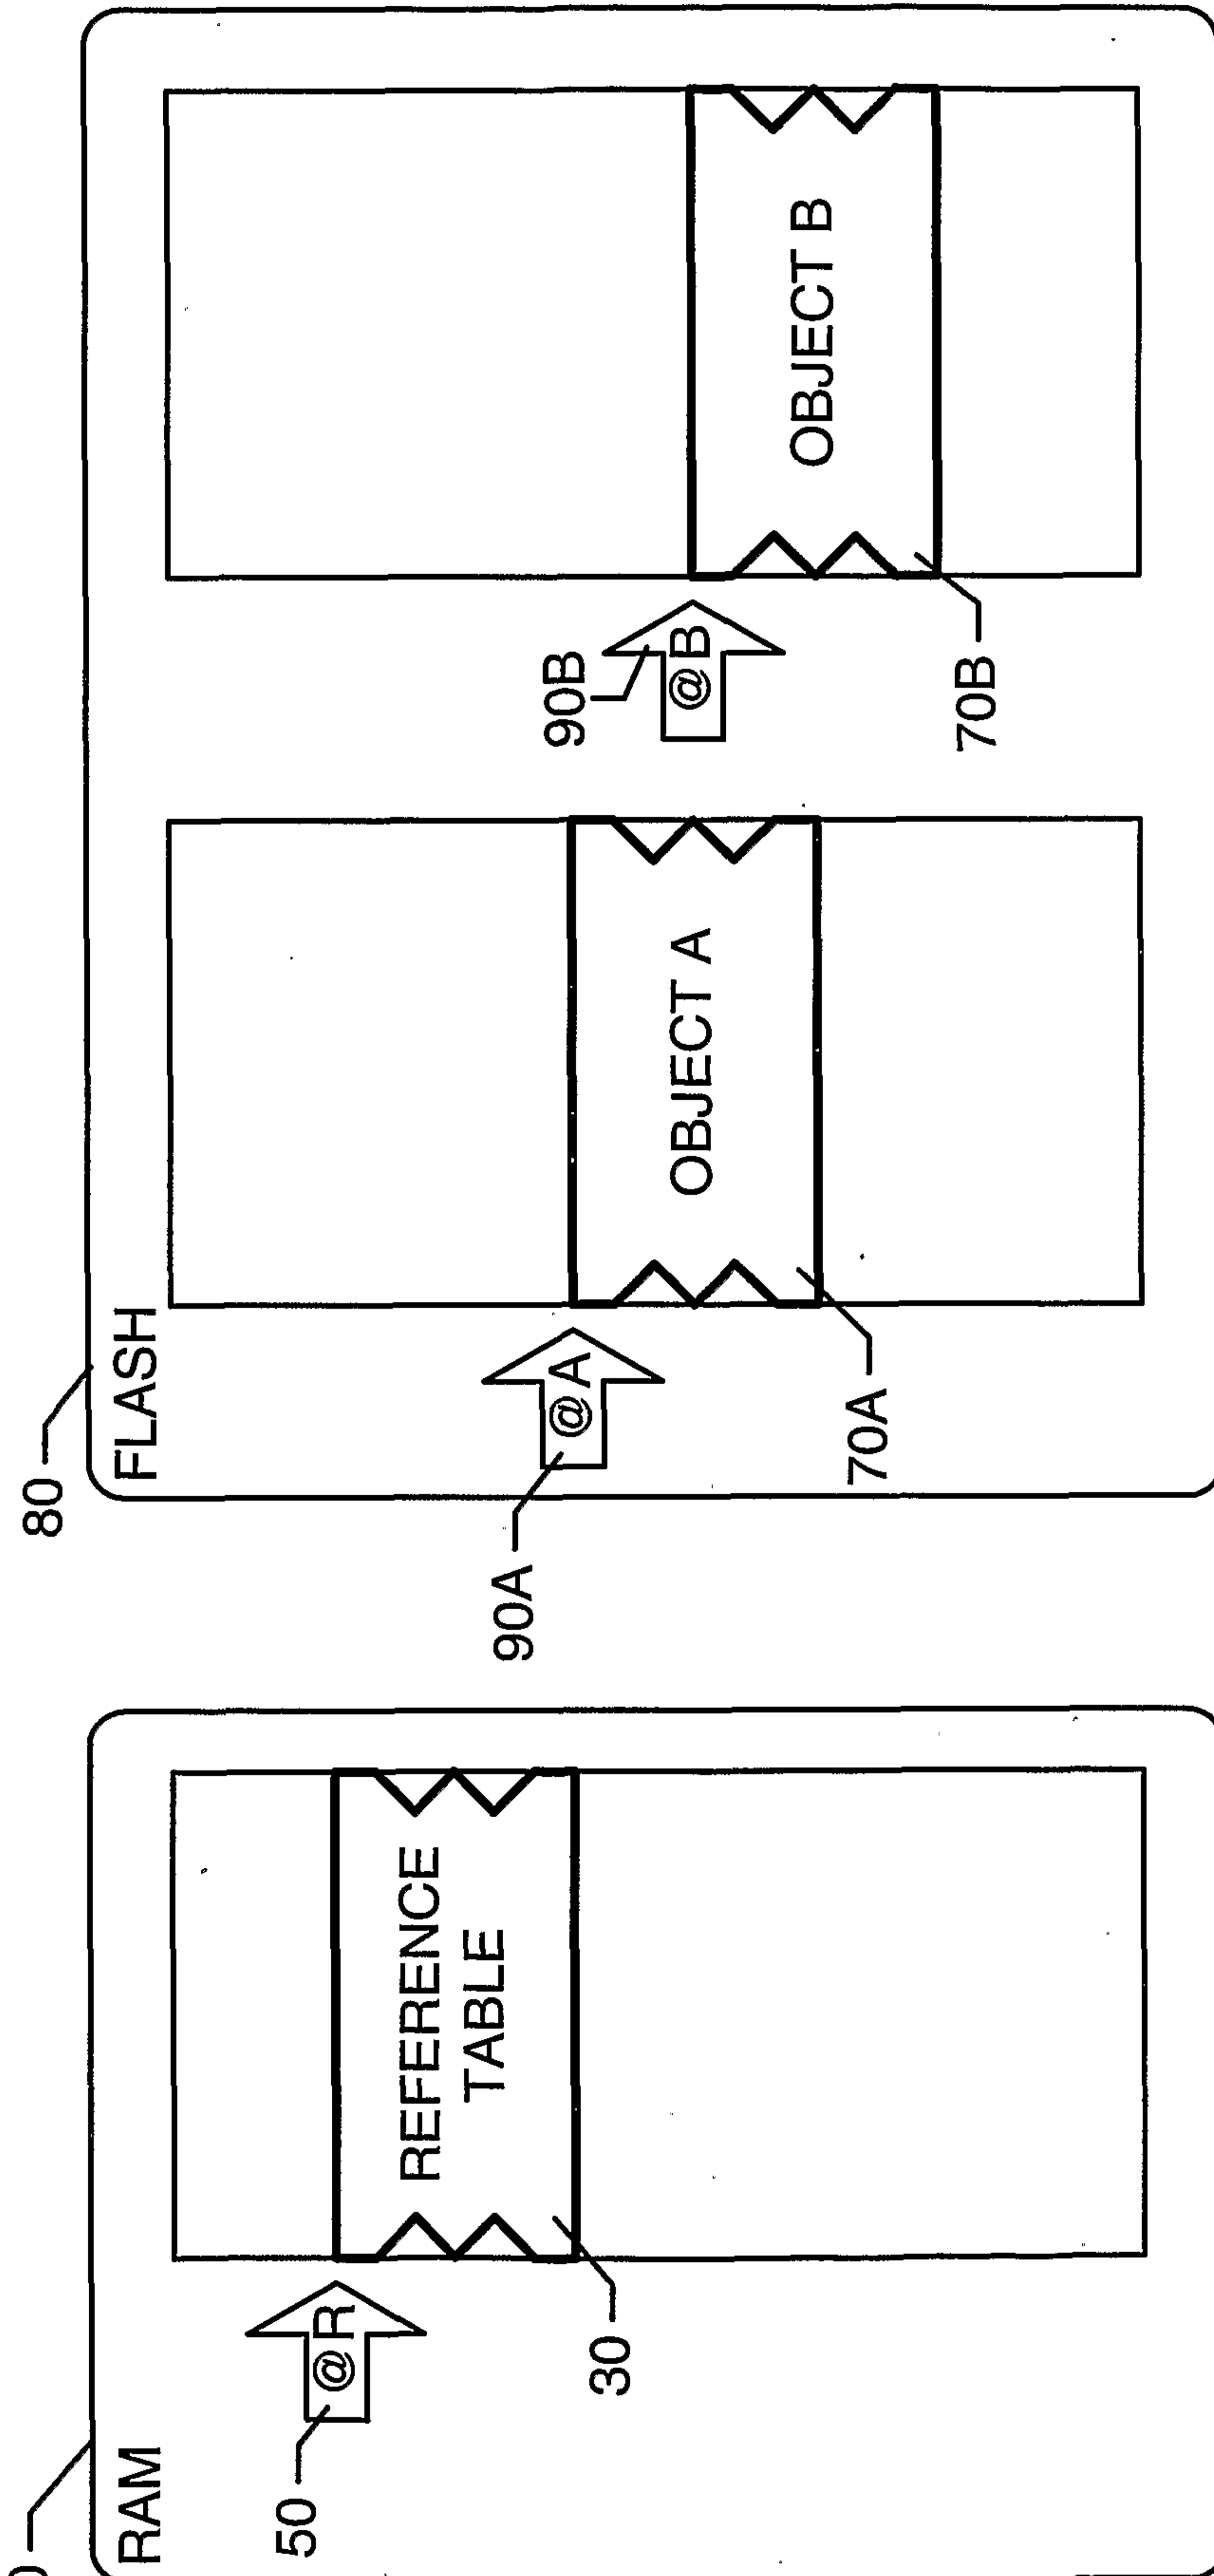


FIG. 5




FIG. 6




FIG. 7




FIG. 8

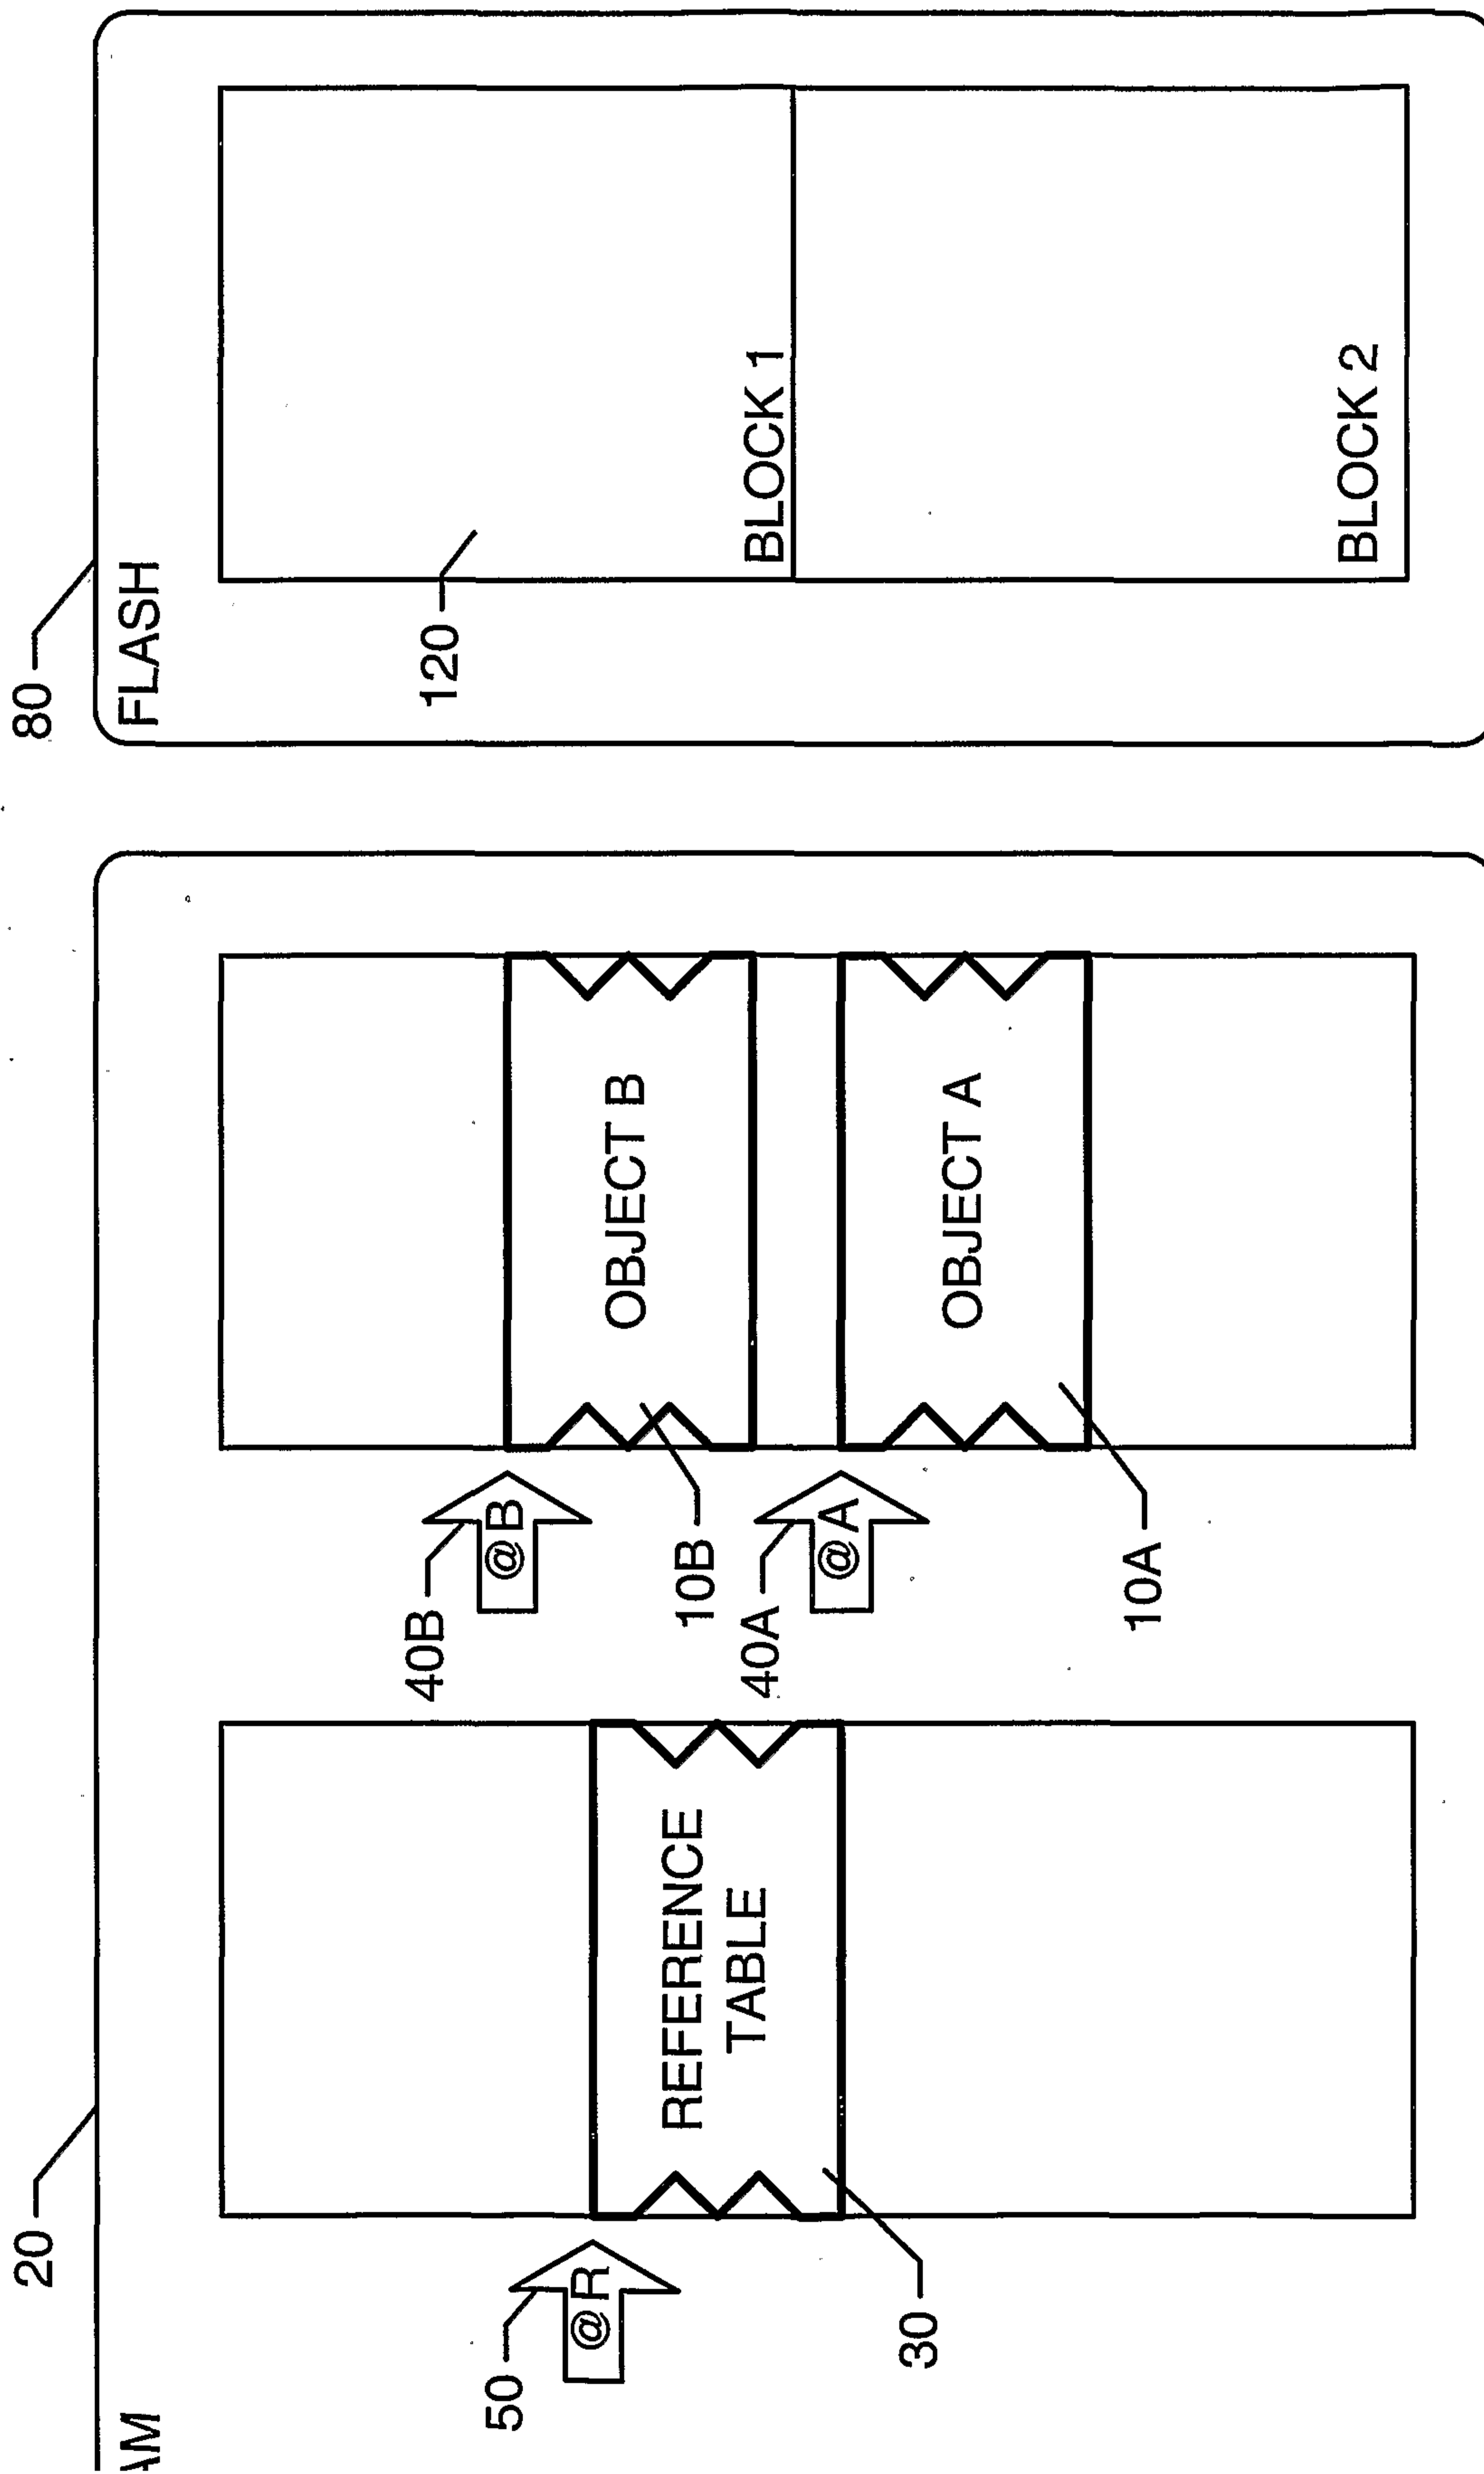


FIG. 9




FIG. 10




FIG. 11 -
GROUPING

12/19




FIG. 12 -group reference table data structure




FIG. 13

n = NUMBER OF
REFERENCE TABLE
ELEMENTS




FIG. 14




FIG. 15

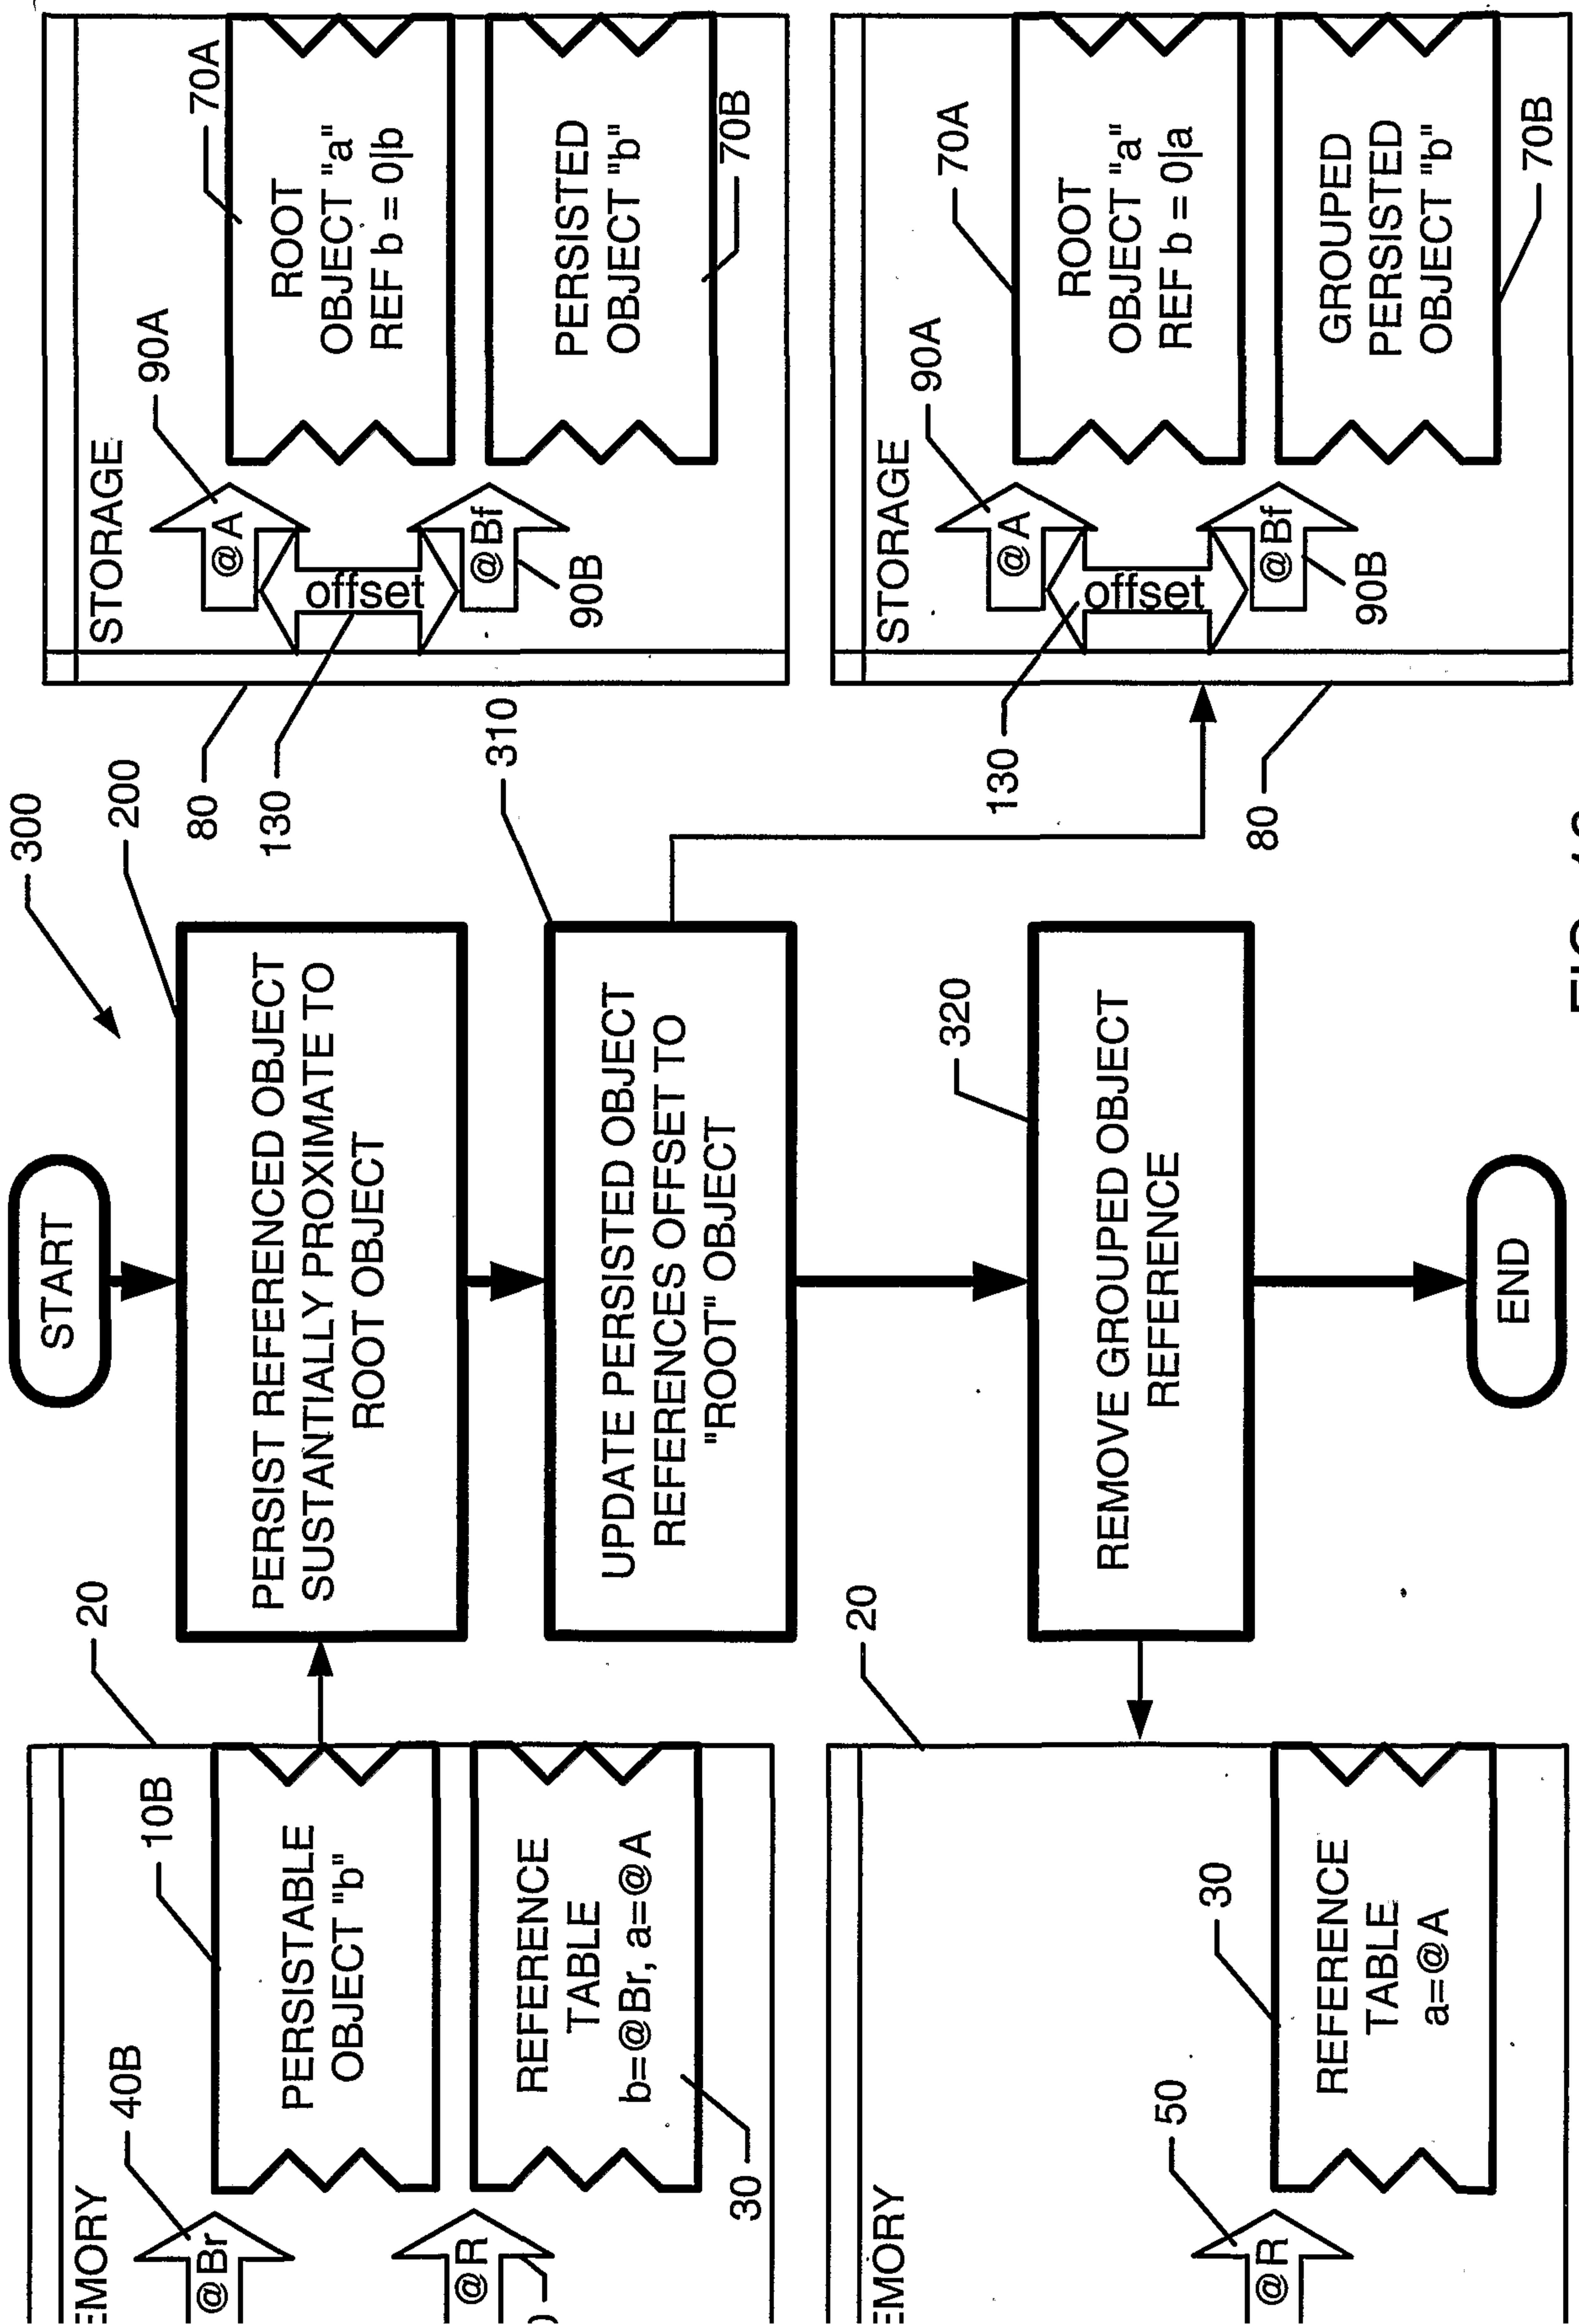


FIG. 16




FIG. 17




FIG. 18




FIG. 19

