

(10) International Publication Number
WO 01/13242 A1

- 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.*

Diagram illustrating a data structure for a card file, showing various fields and their relationships:

- Card Divisions (Left Column):**
 - ONE-CARD DIVISIONS
 - TWO-CARD DIVISIONS
 - * (multiple instances)
 - N-CARD DIVISIONS
 - * (multiple instances)
 - N/2-CARD DIVISIONS
- Main Data Fields (Central Block):**
 - TYPE (108)
 - SECTION SIZE (134)
 - NO_OF_SECTIONS (120)
 - PREVIOUS (122)
 - NEXT (126)
 - FREE SECTIONS (128)
- Object Fields (Below Main Data):**
 - CLASS (112)
 - CNT (114)
 - HASH
 - AGE
 - LK
- Other Object Fields and Car Data (Bottom):**
 - OTHER OBJECT FIELDS
 - CAR DATA
- Card Sequence (Right Column):**
 - A sequence of 'C A R' blocks, representing individual cards.
 - Arrows indicate pointers from the PREVIOUS, NEXT, and FREE SECTIONS fields to specific 'C A R' blocks.
 - A wavy line indicates the continuation of the sequence.

(S7) Abstract: A garbage collector that employs the train algorithm identifies popular objects by a count of references to them and places such objects into special, single-object car sections that are normally smaller than the car sections employed for objects that are less popular. When such cars are part of the collection set and evacuated, the evacuation occurs by relinking, without relocating the objects and memory. References to such popular objects therefore do not need to be updated as a result of evacuation. A limit is imposed on the number of entries in the remembered sets used to keep track of references to objects and such popular-object car sections, and any popular-object car section whose remembered set has more than a threshold number of entries is treated as reachable and still popular without having the memory regions that those entries identify searched for valid references.

POPULAR-OBJECT HANDLING IN A TRAIN- ALGORITHM-BASED GARBAGE COLLECTOR

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is related to commonly assigned U.S. Patent applications of
5 Alexander T. Garthwaite for Scalable-Remembered-Set Garbage Collection, for a
Train-Algorithm-Based Garbage Collector Employing Fixed-Size Remembered Sets,
and for a Train-Algorithm-Based Garbage Collector Employing Reduced Oversized-
Object Threshold, and it is also related to commonly assigned U.S. Patent applications
of Garthwaite et al. for Reduced-Cost Remembered-Set Processing in a Train-
10 Algorithm-Based Garbage Collector and for a Train-Algorithm-Based Garbage Col-
lector Employing Farthest-Forward-Car Indicator, all of which are filed concurrently
herewith and are hereby incorporated in their entirety by reference.

BACKGROUND OF THE INVENTION

The present invention is directed to memory management. It particularly con-
15 cerns what has come to be known as “garbage collection.”

In the field of computer systems, considerable effort has been expended on the
task of allocating memory to data objects. For the purposes of this discussion, the term
object refers to a data structure represented in a computer system’s memory. Other
terms sometimes used for the same concept are *record* and *structure*. An object may be
20 identified by a *reference*, a relatively small amount of information that can be used to
access the object. A reference can be represented as a “pointer” or a “machine ad-
dress,” which may require, for instance, only sixteen, thirty-two, or sixty-four bits of
information, although there are other ways to represent a reference.

In some systems, which are usually known as “object oriented,” objects may
25 have associated *methods*, which are routines that can be invoked by reference to the
object. They also may belong to a *class*, which is an organizational entity that may

contain method code or other information shared by all objects belonging to that class. In the discussion that follows, though, the term *object* will not be limited to such structures; it will additionally include structures with which methods and classes are not associated.

5 The invention to be described below is applicable to systems that allocate memory to objects dynamically. Not all systems employ dynamic allocation. In some computer languages, source programs can be so written that all objects to which the program's variables refer are bound to storage locations at compile time. This storage-allocation approach, sometimes referred to as "static allocation," is the policy tradition-
10 ally used by the Fortran programming language, for example.

 Even for compilers that are thought of as allocating objects only statically, of course, there is often a certain level of abstraction to this binding of objects to storage locations. Consider the typical computer system 10 depicted in Fig. 1, for example. Data, and instructions for operating on them, that a microprocessor 11 uses may reside
15 in on-board cache memory or be received from further cache memory 12, possibly through the mediation of a cache controller 13. That controller 13 can in turn receive such data from system read/write memory ("RAM") 14 through a RAM controller 15 or from various peripheral devices through a system bus 16. The memory space made available to an application program may be "virtual" in the sense that it may actually be
20 considerably larger than RAM 14 provides. So the RAM contents will be swapped to and from a system disk 17.

 Additionally, the actual physical operations performed to access some of the most-recently visited parts of the process's address space often will actually be performed in the cache 12 or in a cache on board microprocessor 11 rather than on the
25 RAM 14, with which those caches swap data and instructions just as RAM 14 and system disk 17 do with each other.

 A further level of abstraction results from the fact that an application will often be run as one of many processes operating concurrently with the support of an underlying operating system. As part of that system's memory management, the application's memory space may be moved among different actual physical locations many
30 times in order to allow different processes to employ shared physical memory devices.

That is, the location specified in the application's machine code may actually result in different physical locations at different times because the operating system adds different offsets to the machine-language-specified location.

Despite these expedients, the use of static memory allocation in writing certain
5 long-lived applications makes it difficult to restrict storage requirements to the available memory space. Abiding by space limitations is easier when the platform provides for dynamic memory allocation, i.e., when memory space to be allocated to a given object is determined only at run time.

Dynamic allocation has a number of advantages, among which is that the run-
10 time system is able to adapt allocation to run-time conditions. For example, the programmer can specify that space should be allocated for a given object only in response to a particular run-time condition. The C-language library function `malloc()` is often used for this purpose. Conversely, the programmer can specify conditions under which memory previously allocated to a given object can be reclaimed for reuse. The C-
15 language library function `free()` results in such memory reclamation.

Because dynamic allocation provides for memory reuse, it facilitates generation of large or long-lived applications, which over the course of their lifetimes may employ objects whose total memory requirements would greatly exceed the available memory resources if they were bound to memory locations statically.

20 Particularly for long-lived applications, though, allocation and reclamation of dynamic memory must be performed carefully. If the application fails to reclaim unused memory—or, worse, loses track of the address of a dynamically allocated segment of memory—its memory requirements will grow over time to exceed the system's available memory. This kind of error is known as a “memory leak.”

25 Another kind of error occurs when an application reclaims memory for reuse even though it still maintains a reference to that memory. If the reclaimed memory is reallocated for a different purpose, the application may inadvertently manipulate the same memory in multiple inconsistent ways. This kind of error is known as a “dangling reference,” because an application should not retain a reference to a memory lo-

cation once that location is reclaimed. Explicit dynamic-memory management by using interfaces like malloc()/free() often leads to these problems.

A way of reducing the likelihood of such leaks and related errors is to provide memory-space reclamation in a more-automatic manner. Techniques used by systems that reclaim memory space automatically are commonly referred to as “garbage collection.” Garbage collectors operate by reclaiming space that they no longer consider “reachable.” Statically allocated objects represented by a program’s global variables are normally considered reachable throughout a program’s life. Such objects are not ordinarily stored in the garbage collector’s managed memory space, but they may contain references to dynamically allocated objects that are, and such objects are considered reachable. Clearly, an object referred to in the processor’s call stack is reachable, as is an object referred to by register contents. And an object referred to by any reachable object is also reachable.

The use of garbage collectors is advantageous because, whereas a programmer working on a particular sequence of code can perform his task creditably in most respects with only local knowledge of the application at any given time, memory allocation and reclamation require a global knowledge of the program. Specifically, a programmer dealing with a given sequence of code does tend to know whether some portion of memory is still in use for that sequence of code, but it is considerably more difficult for him to know what the rest of the application is doing with that memory. By tracing references from some conservative notion of a “root set,” e.g., global variables, registers, and the call stack, automatic garbage collectors obtain global knowledge in a methodical way. By using a garbage collector, the programmer is relieved of the need to worry about the application’s global state and can concentrate on local-state issues, which are more manageable. The result is applications that are more robust, having no dangling references and fewer memory leaks.

Garbage-collection mechanisms can be implemented by various parts and levels of a computing system. One approach is simply to provide them as part of a batch compiler’s output. Consider Fig. 2’s simple batch-compiler operation, for example. A computer system executes in accordance with compiler object code and therefore acts as a compiler. The compiler object code is typically stored on a medium such as

- 5 -

Fig. 1's system disk 17 or some other machine-readable medium, and it is loaded into RAM 14 to configure the computer system to act as a compiler. In some cases, though, the compiler object code's persistent storage may instead be provided in a server system remote from the machine that performs the compiling. The electrical signals that carry the digital data by which the computer systems exchange that code are exemplary forms of carrier waves transporting the information.

The input to the compiler is the application source code, and the end product of the compiler process is application object code. This object code defines an application 21, which typically operates on input such as mouse clicks, etc., to generate a display or some other type of output. This object code implements the relationship that the programmer intends to specify by his application source code. In one approach to garbage collection, the compiler 20, without the programmer's explicit direction, additionally generates code that automatically reclaims unreachable memory space.

Even in this simple case, though, there is a sense in which the application does not itself provide the entire garbage collector. Specifically, the application will typically call upon the underlying operating system's memory-allocation functions. And the operating system may in turn take advantage of various hardware that lends itself particularly to use in garbage collection. So even a very simple system may disperse the garbage-collection mechanism over a number of computer-system layers.

To get some sense of the variety of system components that can be used to implement garbage collection, consider Fig. 3's example of a more complex way in which various levels of source code can result in the machine instructions that a processor executes. In the Fig. 3 arrangement, the human applications programmer produces source code 22 written in a high-level language. A compiler 23 typically converts that code into "class files." These files include routines written in instructions, called "byte codes" 24, for a "virtual machine" that various processors can be configured to emulate. This conversion into byte codes is almost always separated in time from those codes' execution, so Fig. 3 divides the sequence into a "compile-time environment" 25 separate from a "run-time environment" 26, in which execution occurs. One example of a high-level language for which compilers are available to produce such virtual-machine instructions is the Java™ programming language. (Java is a trademark or

registered trademark of Sun Microsystems, Inc., in the United States and other countries.)

Most typically, the class files' byte-code routines are executed by a processor under control of a virtual-machine process 27. That process emulates a virtual machine from whose instruction set the byte codes are drawn. As is true of the compiler 23, the
5 virtual-machine process 27 may be specified by code stored on a local disk or some other machine-readable medium from which it is read into Fig. 1's RAM 14 to configure the computer system to act as a virtual machine and implement the garbage collector. Again, though, that code's persistent storage may instead be provided by a server
10 system remote from the processor that implements the virtual machine, in which case the code would be transmitted electrically or optically to the virtual-machine-implementing processor.

In most implementations, much of the virtual machine's action in executing these byte codes is most like what those skilled in the art refer to as "interpreting," and
15 Fig. 3 shows that the virtual machine includes an "interpreter" 28 for that purpose. The resultant instructions typically invoke calls to a run-time system 29, which handles matters such as loading new class files as they are needed and, of particular interest in the present connection, performing garbage collection.

In addition to running an interpreter, many virtual-machine implementations
20 also actually compile the byte codes concurrently with the resultant object code's execution, so Fig. 3 depicts the virtual machine as additionally including a "just-in-time" compiler 30.

The arrangement of Fig. 3 differs from Fig. 2 in that the compiler 23 for converting the human programmer's code does not contribute to providing the garbage-
25 collection function; that results largely from the virtual machine 27's operation. Although the Fig. 3 arrangement is a popular one, it is by no means universal, and many further implementation types can be expected. Proposals have even been made to implement the virtual machine 27's behavior in a hardware processor, in which case the hardware itself would provide some or all of the garbage-collection function.

In short, garbage collectors can be implemented in a wide range of combinations of hardware and/or software. As is true of most of the garbage-collection techniques described in the literature, the invention to be described below is applicable to most such systems.

5 By implementing garbage collection, a computer system can greatly reduce the occurrence of memory leaks and other software deficiencies in which human programming frequently results. But it can also have significant adverse performance effects if it is not implemented carefully. To distinguish the part of the program that does “useful” work from that which does the garbage collection, the term *mutator* is sometimes
10 used in discussions of these effects; from the collector’s point of view, what the mutator does is mutate active data structures’ connectivity.

Some garbage-collection approaches rely heavily on interleaving garbage-collection steps among mutator steps. In one type of garbage-collection approach, for instance, the mutator operation of writing a reference is followed immediately by garbage-collector steps used to maintain a reference count in that object’s header, and code
15 for subsequent new-object storage includes steps for finding space occupied by objects whose reference count has fallen to zero. Obviously, such an approach can slow mutator operation significantly.

Other approaches therefore interleave very few garbage-collector-related instructions into the main mutator process but instead interrupt it from time to time to
20 perform garbage-collection cycles, in which the garbage collector finds unreachable objects and reclaims their memory space for reuse. Such an approach will be assumed in discussing Fig. 4’s depiction of a simple garbage-collection operation. Within the memory space allocated to a given application is a part 40 managed by automatic garbage collection. In the following discussion, this will be referred to as the “heap,” although in other contexts that term refers to all dynamically allocated memory. During
25 the course of the application’s execution, space is allocated for various objects 42, 44, 46, 48, and 50. Typically, the mutator allocates space within the heap by invoking the garbage collector, which at some level manages access to the heap. Basically, the mutator asks the garbage collector for a pointer to a heap region where it can safely place
30 the object’s data. The garbage collector keeps track of the fact that the thus-allocated

- 8 -

region is occupied. It will refrain from allocating that region in response to any other request until it determines that the mutator no longer needs the region allocated to that object.

Garbage collectors vary as to which objects they consider reachable and un-
5 reachable. For the present discussion, though, an object will be considered “reachable” if it is referred to as object 42 is, by a reference in the root set 52. The root set consists of reference values stored in the mutator’s threads’ call stacks, the CPU registers, and global variables outside the garbage-collected heap. An object is also reachable if it is referred to, as object 46 is, by another reachable object (in this case, object 42). Ob-
10 jects that are not reachable can no longer affect the program, so it is safe to re-allocate the memory spaces that they occupy.

A typical approach to garbage collection is therefore to identify all reachable objects and reclaim any previously allocated memory that the reachable objects do not occupy. A typical garbage collector may identify reachable objects by tracing refer-
15 ences from the root set 52. For the sake of simplicity, Fig. 4 depicts only one reference from the root set 52 into the heap 40. (Those skilled in the art will recognize that there are many ways to identify references, or at least data contents that may be references.) The collector notes that the root set points to object 42, which is therefore reachable, and that reachable object 42 points to object 46, which therefore is also reachable. But
20 those reachable objects point to no other objects, so objects 44, 48, and 50 are all unreachable, and their memory space may be reclaimed.

To avoid excessive heap fragmentation, some garbage collectors additionally relocate reachable objects. Fig. 5 shows a typical approach. The heap is partitioned into two halves, hereafter called “semi-spaces.” For one garbage-collection cycle, all
25 objects are allocated in one semi-space 54, leaving the other semi-space 56 free. When the garbage-collection cycle occurs, objects identified as reachable are “evacuated” to the other semi-space 56, so all of semi-space 54 is then considered free. Once the garbage-collection cycle has occurred, all new objects are allocated in the lower semi-space 56 until yet another garbage-collection cycle occurs, at which time the reachable
30 objects are evacuated back to the upper semi-space 54.

- 9 -

Although this relocation requires the extra steps of copying the reachable objects and updating references to them, it tends to be quite efficient, since most new objects quickly become unreachable, so most of the current semi-space is actually garbage. That is, only a relatively few, reachable objects need to be relocated, after which
5 the entire semi-space contains only garbage and can be pronounced free for reallocation.

In one sense, the approach of interrupting the mutator occasionally for garbage collection can increase an application's responsiveness, because the main mutator operation ordinarily proceeds relatively unburdened by garbage-collection overhead. In
10 interactive systems, moreover, interruptions for garbage collection can sometimes be scheduled opportunistically so as to reduce the likelihood that they will result in much overall speed reduction. Garbage collection can be triggered when the system is waiting for user input, for instance.

So it may often be true that the garbage-collection operation's effect on performance can depend less on the total collection time than on when collections actually
15 occur. But another factor that often is even more determinative is the duration of any single collection cycle, i.e., how long the mutator must remain quiescent at any one time. In an interactive system, for instance, a user may never notice hundred-millisecond interruptions for garbage collection, whereas most users would find interruptions lasting for two seconds to be annoying. Many garbage collectors therefore operate
20 incrementally. That is, they perform less than a complete collection in any single interruption of the main application.

A way of not only reducing collection-cycle length but also increasing overall efficiency is to segregate the heap into one or more parts, called generations, that are
25 subject to different collection policies. New objects are allocated in a "young" generation, and older objects are promoted from younger generations to older or more "mature" generations. Collecting the younger generations more frequently than the others yields greater efficiency because the younger generations tend to accumulate garbage faster; newly allocated objects tend to "die," while older objects tend to "survive."

30 But generational collection greatly increases what is effectively the root set for a given generation. Consider Fig. 6, which depicts a heap as organized into three gen-

- 10 -

erations 58, 60, and 62. Assume that generation 60 is to be collected. The process for this individual generation may be more or less the same as that described in connection with Figs. 4 and 5 for the entire heap, with one major exception. In the case of a single generation, the root set must be considered to include not only the call stack, registers, and global variables represented by set 52 but also objects in the other generations 58 and 62, which themselves may contain references to objects in generation 60. So pointers must be traced not only from the basic root set 52 but also from objects within the other generations.

One could perform this tracing by simply inspecting all references in all other generations at the beginning of every collection cycle, and it turns out that this approach is actually feasible in some situations. But it takes too long in other situations, so workers in this field have employed a number of approaches to expediting reference tracing. One approach is to include so-called write barriers in the mutator process. A *write barrier* is code added to a write operation to record information from which the collector can determine where references were or may have been written since the last collection cycle. A reference list can then be maintained by taking such a list as it existed at the end of the previous collection cycle and updating it by inspecting only locations identified by the write barrier as possibly modified since the last collection cycle.

One of the many such implementations commonly used by workers in this art employs what has been referred to as the "card table." Fig. 6 depicts the various generations as being divided into smaller sections, known for this purpose as "cards." Card tables 64, 66, and 68 associated with respective generations contain an entry for each of their cards. Each card-table entry includes some summary of its associated card's inter-generational-reference content. This summary may be no more than a binary indication of whether such a reference exists, but it preferably includes a list of offsets that indicate such references' locations in the card. Part of entry 69, for instance, may be an offset that indicates the location of an inter-generational reference 70 within the corresponding card 72. For the sake of concreteness, we will assume that the summaries are maintained by steps that occur principally at the beginning of each collection cycle.

Of course, the card-table approach is only one of many that can be employed to detect inter-generational pointers. Indeed, it is typical for an individual garbage col-

- 11 -

lector to use more than one approach. Although there is no reason in principle to favor any particular number of generations, and although Fig. 6 shows three, most generational garbage collectors have only two generations, of which one is the young generation and the other is the mature generation. Moreover, although Fig. 6 shows the generations as being of the same size, a more-typical configuration is for the young generation to be considerably smaller. Finally, although we assumed for the sake of simplicity that collection during a given cycle was limited to only one generation, a more-typical approach is actually to collect the whole young generation at every cycle but to collect the mature one less frequently.

10 To collect the young generation, it is preferable to employ the card table to identify pointers into the young generation; laboriously scanning the entire mature generation would take too long. On the other hand, since the young generation is collected in every cycle and can therefore be collected before mature-generation processing, it takes little time to scan the few remaining, live objects in the young generation for
15 pointers into the mature generation in order to process that generation. For this reason, the card table will typically be so maintained as only to identify the regions occupied by references into younger generations and not into older ones.

Now, although it typically takes very little time to collect the young generation, it may take more time than is acceptable within a single garbage-collection cycle to
20 collect the entire mature generation. So some garbage collectors may collect the mature generation incrementally; that is, they may perform only a part of the mature generation's collection during any particular collection cycle. Incremental collection presents the problem that, since the generation's objects that are outside a collection cycle's collection set are not processed during that cycle, any such objects that are unreachable
25 are not recognized as unreachable, so collection-set objects to which they refer tend not to be, either.

To reduce the adverse effect this would otherwise have on collection efficiency, workers in this field have employed the "train algorithm," which Fig. 7 depicts. A generation to be collected incrementally is divided into sections, which for reasons about to
30 be described are referred to as "car sections." The generation may be collected incrementally in fixed-size sections, in which case the size may equal that of the generation

portion to be collected during one cycle. More typically, the fraction of the generation to be collected during any one cycle varies with run-time conditions, in which case the car-section size is often that of the minimum fraction to be collected in one cycle, although that is not a requirement.

5 Additionally, the car sections are grouped into “trains,” which are ordered according to age. For example, Fig. 7 shows an oldest train 73 consisting of a generation 74’s three car sections described by associated data structures 75, 76, and 78, while a second train 80 consists only of a single car section, represented by structure 82, and the youngest train 84 (referred to as the “allocation train”) consists of car sections that
10 data structures 86 and 88 represent. As will be seen below, car sections’ train memberships can change, and any car section added to a train is added to the end of that train. Train size is a matter of design choice, but its purpose is to maximize the probability that garbage reference “cycles” can be reclaimed, as will now be explained.

15 Although more than one car section may be collected in any given cycle, little generality is lost by considering the situation in which only a single car is collected during any cycle. The car collected in any cycle is the one added least recently to the oldest train, which in this case is car 75. All of the generation’s cars can thus be thought of as waiting for collection in a single long line, in which cars in a given train are farther forward than those in any younger train, and those added to a train more re-
20 cently are behind those added to it earlier.

As is usual, the way in which reachable objects are identified is to determine whether there are references to them in the root set or in any other object already determined to be reachable. In accordance with the train algorithm, the collector additionally performs a test to determine whether there are any references at all from out-
25 side the oldest train to objects within it. If there are not, then all cars within the train can be reclaimed, even though not all of those cars are in the collection set. And the train algorithm so operates that inter-car references tend to be grouped into trains, as will now be explained.

To identify references into the car from outside of it, train-algorithm imple-
30 mentations typically employ “remembered sets.” As card tables are, remembered sets are used to keep track of references. Whereas a card-table entry contains information

about references that the associated card contains, though, a remembered set associated with a given region contains information about references into that region from locations outside of it. In the case of the train algorithm, remembered sets are associated with car sections. Each remembered set, such as car 75's remembered set 90, lists locations in the generation that contain references into the associated car section. The remembered sets for all of a generation's cars are typically updated at the start of each collection cycle, concurrently with card-table updates. For reasons that will become apparent, the collector will typically not bother to place in the remembered set the locations of references from objects in car sections farther forward in the collection queue, i.e., from objects in older trains or cars added earlier to the same train. For the sake of simplicity, we will continue the assumption that only a single car is collected during each collection cycle, although we will discuss multiple-car collection sets presently.

When the remembered sets have been updated, the collector reads the remembered set of each car in the collection set to determine the location of each reference from a higher-order car into the collection set. The collector places the address of each reference thereby found into a scratch-pad list associated with the train that contains that reference. (Again, the remembered set lists only reference locations in the same generation).

When the collector has read all references in the remembered set, it evacuates into the youngest train the collection-set-car objects referred to by the references in the locations that the youngest train's scratch-pad-list entries specify. It also removes those scratch-pad-list entries and updates the references to which they pointed so that those references reflect the evacuated objects' new locations. Any collection-set objects to which the thus-evacuated objects refer are similarly evacuated to that train, and this continues until that train no longer contains any references into the collection-set car section.

Whenever an object is evacuated, the collector leaves an indication of this fact in the object's previous location, together with the address of its new location. So, if the reference found in the location identified by any subsequent scratch-pad-list entry

refers to an already-evacuated object, the collector is apprised of this fact and can update the reference without attempting to evacuate the already-evacuated object.

This process is repeated for successively older trains until the collector reaches the oldest train. Before it processes references in that train's scratch-pad list, the collector evacuates any objects referred to from outside the generation. To identify such
5 objects, the collector scans the root set and other generations for references into the collection set. Now, it may not be necessary to scan all other generations. A particularly common scheme is not to collect any generation in a collection cycle in which every younger generation is not completely collected, and the policy may be to promote
10 all surviving younger-generation objects into older generations. In such a situation, it is necessary only to scan older generations.

The scanning may actually involve inspecting each surviving object in the other generation, or the collector may expedite the process by using card-table entries. Regardless of which approach it uses, the collector immediately evacuates into another
15 train any collection-set object to which it thereby finds an external reference. The typical policy is to place the evacuated object into the youngest such train. As before, the collector does not attempt to evacuate an object that has already been evacuated, and, when it does evacuate an object to a train, it evacuates to the same train any objects in the collection-set car to which the thus-evacuated object refers. In any case, the collector
20 updates the reference to the evacuated object.

When all inter-generationally referred-to objects have been evacuated from the collection-set car, the collector proceeds to evacuate any collection-set objects referred to by references whose locations the oldest train's scratch-pad list includes. It removes them to younger cars in the oldest train, again updating references, avoiding duplicate
25 evacuations, and evacuating any collection-set-car objects to which the evacuated objects refer. When this process has been completed, the car section can be reclaimed, since any reference to any remaining object must reside in the same car, so all remaining collection-set objects are unreachable.

When the collection-set car section has been reclaimed, the garbage collector
30 then performs the train algorithm's central test: it determines whether there are any references into the oldest train from outside that train. If not, the entire train can be re-

- 15 -

claimed, even if there are inter-car references between its individual cars. By evacuat-
ing objects into the cars that references to them occupy, the train algorithm tends to
group garbage reference cycles into single trains, whose sizes are not limited, as car
sizes are, by the need to optimize collection-cycle duration. The train algorithm is thus
5 able to reclaim space occupied by large garbage reference cycles even if the space in-
crements (car sections) that it collects are relatively small. To support this process, a
tally of how many references there are from other trains in the same generation is typi-
cally maintained in connection with the various remembered-set updates. This tally,
together with a tally of extra-generational references developed during the extra-
10 generational scan, yields the indications of whether there are any references into a
given train from outside that train.

Figs. 8A-8J illustrate results of using the train algorithm. Fig. 8A represents a
generation in which objects have been allocated in nine car sections. The oldest train
has four cars, numbered 1.1 through 1.4. Car 1.1 has two objects, A and B. There is a
15 reference to object B in the root set (which, as was explained above, includes live ob-
jects in the other generations). Object A is referred to by object L, which is in the third
train's sole car section. In the generation's remembered sets 92, a reference in object L
has therefore been recorded against car 1.1.

Processing always starts with the oldest train's earliest-added car, so the gar-
20 bage collector refers to car 1.1's remembered set and finds that there is a reference from
object L into the car being processed. It accordingly evacuates object A to the train that
object L occupies. The object being evacuated is often placed in one of the selected
train's existing cars, but we will assume for present purposes that there is not enough
room. So the garbage collector evacuates object A into a new car section and updates
25 appropriate data structures to identify it as the next car in the third train. Fig. 8B de-
picts the result: a new car has been added to the third train, and object A is placed in it.

Fig. 8B also shows that object B has been evacuated to a new car outside the
first train. This is because object B has an external reference, which, like the reference
to object A, is a reference from outside the first train, and one goal of the processing is
30 to form trains into which there are no further references. Note that, to maintain a refer-
ence to the same object, object L's reference to object A has had to be rewritten, and so

- 16 -

have object B's reference to object A and the inter-generational pointer to object B. In the illustrated example, the garbage collector begins a new train for the car into which object B is evacuated, but this is not a necessary requirement of the train algorithm. That algorithm requires only that externally referenced objects be evacuated to a newer
5 train.

Since car 1.1 no longer contains live objects, it can be reclaimed, as Fig. 8B also indicates. Also note that the remembered set for car 2.1 now includes the address of a reference in object A, whereas it did not before. As was stated before, remembered sets in the illustrated embodiment include only references from cars further back in the or-
10 der than the one with which the remembered set is associated. The reason for this is that any other cars will already be reclaimed by the time the car associated with that remembered set is processed, so there is no reason to keep track of references from them.

The next step is to process the next car, the one whose index is 1.2. The pur-
15 pose of a generation's division into car sections is to provide for incremental collection, so collection of car 1.2 can wait until the collection cycle following the one during which car 1.1 is collected. Indeed, the garbage collector can postpone its processing for several collection cycles (during which presumably only younger generations are col-
lected). As was mentioned above, we will assume in this scenario that each car's col-
20 lection occurs in its own separate collection cycle. And for the sake of simplicity we will assume that the mutator does not change any references into the generation in the interim.

Fig. 8B depicts car 1.2 as containing only a single object, object C, and that car's remembered set contains the address of an inter-car reference from object F. The
25 garbage collector follows that reference to object C. Since this identifies object C as possibly reachable, the garbage collector evacuates it from car set 1.2, which is to be reclaimed. Specifically, the garbage collector removes object C to a new car section, section 1.5, which is linked to the train to which the referring object F's car belongs. Of course, object F's reference needs to be updated to object C's new location. Fig. 8C
30 depicts the evacuation's result.

- 17 -

Fig. 8C also indicates that car set 1.2 has been reclaimed, and car 1.3 is next to be processed. The only address in car 1.3's remembered set is that of a reference in object G. Inspection of that reference reveals that it refers to object F. Object F may therefore be reachable, so it must be evacuated before car section 1.3 is reclaimed. On the other hand, there are no references to objects D and E, so they are clearly garbage. Fig. 8D depicts the result of reclaiming car 1.3's space after evacuating possibly reachable object F.

In the state that Fig. 8D depicts, car 1.4 is next to be processed, and its remembered set contains the addresses of references in objects K and C. Inspection of object K's reference reveals that it refers to object H, so object H must be evacuated. Inspection of the other remembered-set entry, the reference in object C, reveals that it refers to object G, so that object is evacuated, too. As Fig. 8E illustrates, object H must be added to the second train, to which its referring object K belongs. In this case there is room enough in car 2.2, which its referring object K occupies, so evacuation of object H does not require that object K's reference to object H be added to car 2.2's remembered set. Object G is evacuated to a new car in the same train, since that train is where referring object C resides. And the address of the reference in object G to object C is added to car 1.5's remembered set.

Fig. 8E shows that this processing has eliminated all references into the first train, and it is an important part of the train algorithm to test for this condition. That is, even though there are references into both of the train's cars, those cars' contents can be recognized as all garbage because there are no references into the train from outside of it. So all of the first train's cars are reclaimed.

The collector accordingly processes car 2.1 during the next collection cycle, and that car's remembered set indicates that there are two references outside the car that refer to objects within it. Those references are in object K, which is in the same train, and object A, which is not. Inspection of those references reveals that they refer to objects I and J, which are evacuated.

The result, depicted in Fig. 8F, is that the remembered sets for the cars in the second train reveal no inter-car references, and there are no inter-generational refer-

ences into it, either. That train's car sections therefore contain only garbage, and their memory space can be reclaimed.

So car 3.1 is processed next. Its sole object, object L, is referred to inter-generationally as well as by a reference in the fourth train's object M. As Fig. 8G shows, object L is therefore evacuated to the fourth train. And the address of the reference in object L to object A is placed in the remembered set associated with car 3.2, in which object A resides.

The next car to be processed is car 3.2, whose remembered set includes the addresses of references into it from objects B and L. Inspection of the reference from object B reveals that it refers to object A, which must therefore be evacuated to the fifth train before car 3.2 can be reclaimed. Also, we assume that object A cannot fit in car section 5.1, so a new car 5.2 is added to that train, as Fig. 8H shows, and object A is placed in its car section. All referred-to objects in the third train having been evacuated, that (single-car) train can be reclaimed in its entirety.

A further observation needs to be made before we leave Fig. 8G. Car 3.2's remembered set additionally lists a reference in object L, so the garbage collector inspects that reference and finds that it points to the location previously occupied by object A. This brings up a feature of copying-collection techniques such as the typical train-algorithm implementation. When the garbage collector evacuates an object from a car section, it marks the location as having been evacuated and leaves the address of the object's new location. So when the garbage collector traces the reference from object L, it finds that object A has been removed, and it accordingly copies the new location into object L as the new value of its reference to object A.

In the state that Fig. 8H illustrates, car 4.1 is the next to be processed. Inspection of the fourth train's remembered sets reveals no inter-train references into it, but the inter-generational scan (possibly performed with the aid of Fig. 6's card tables) reveals inter-generational references into car 4.2. So the fourth train cannot be reclaimed yet. The garbage collector accordingly evacuates car 4.1's referred-to objects in the normal manner, with the result that Fig. 8I depicts.

In that state, the next car to be processed has only inter-generational references into it. So, although its referred-to objects must therefore be evacuated from the train, they can be placed in any other train. In the illustrated implementation, a new train is formed for this purpose, so the result of car 4.2's processing is the state that Fig. 8J depicts.

Processing continues in this same fashion. Of course, subsequent collection cycles will not in general proceed, as in the illustrated cycles, without any reference changes by the mutator and without any addition of further objects. But reflection reveals that the general approach just described still applies when such mutations occur.

However, there is a simplification in the foregoing scenario that does obscure a difficulty encountered in implementing the train algorithm. Fig. 8F shows that there are two references to object L after the second train is collected. So references in both of the referring objects need to be updated when object L is evacuated. This is not a great burden, since only two referring objects are involved, but some types of applications routinely generate objects that are very "popular," i.e., are objects to which there are large numbers of references. Evacuating a single one of these objects therefore requires considerable reference updating.

If the object remains popular, moreover, the number of entries that must thereafter be made to its car's remembered set during later collection cycles will also be high. This causes such remembered sets to become large and unwieldy. In the worst case, such an object can be referred to by almost every object in the generation, in which case the remembered set would be on the order of the entire generation's size. This not only causes a significant space problem but also makes maintaining and processing remembered sets costly. As a remembered set's size increases, the cost of adding new entries, eliminating duplicate entries for the same references, and scanning the references during the car's collection can become unacceptable. The garbage-collection overhead thereby imposed by popular-object-using applications may be so great as to make it impractical to provide garbage-collection intervals that are short enough to meet performance requirements.

One proposal for dealing with popular objects involves marking such objects' cars as popular and simply exempting those cars from collection, on the theory that

popular objects rarely become unreachable. This eliminates the need to maintain remembered sets for those cars, so there are no oversized remembered sets. But one can readily conceive of pathological cases in which large numbers of cars contain such objects that have in fact died. The heap could become largely useless in such cases.

5 Another approach is to take advantage of certain features of object-oriented languages. It may be determined from experience that the few objects that become popular are all instances of a small number of classes. So objects that are instances of those classes can be allocated in a special heap not managed in accordance with the train algorithm. That, too, eliminates the remembered-set problem that they would otherwise
10 present. But the classes whose instances tend to be popular are different for different kinds of applications, and obtaining such knowledge for different applications would be impractical.

SUMMARY OF THE INVENTION

 The present invention provides a way of avoiding many popular-object problems without exempting them from collection and without requiring advance knowl-
15 edge of which classes tend to have popular instances. According to the present invention, special car sections are set aside to contain popular objects. If an object in the collection set in any given collection cycle meets some pre-determined popularity indication, typically a threshold number of references to it, then it is placed in one such
20 popular-object car section. When a popular-object car section contains one object, the collector places no more objects in that car section. Popular-object car sections that thus contain popular objects become members of trains, just as normal car sections do, and they similarly take their turn at being collected. However, the fact that a popular-object car section contains only one object enables the collection operation to be per-
25 formed by simply re-linking the popular-object car section, without relocating its contained popular object. As a consequence, references to the popular object do not have to be rewritten when the popular-object car section is collected.

 Preferably, popular-object car sections are provided in many sizes, most of which typically will be much smaller than the ordinary car section's size, which is
30 normal chosen to match the granularity of collection-set variation and is much larger than the typical object.

This approach can be implemented quite flexibly. In particular, it can not only test objects in normal car sections to determine whether they are popular and should be transferred to popular-car sections, but some embodiments will further test objects in popular-car sections to determine whether they have ceased to be popular and should therefore be placed back into normal car sections.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention description below refers to the accompanying drawings, of which:

Fig. 1 is a block diagram of a computer system in which the present invention's teachings can be practiced;

Fig. 2 is a block diagram that illustrates a compiler's basic functions;

Fig. 3 is a block diagram that illustrates a more-complicated compiler/interpreter organization;

Fig. 4 is a diagram that illustrates a basic garbage-collection mechanism;

Fig. 5 is a similar diagram illustrating that garbage-collection approach's relocation operation;

Fig. 6 is a diagram that illustrates a garbage-collected heap's organization into generations;

Fig. 7 is a diagram that illustrates a generation organization employed for the train algorithm;

Figs. 8A-J are diagrams illustrating a collection scenario that can result from use of the train algorithm;

Fig. 9 is a diagram that illustrates train and car data structures that may be used in support of the train algorithm;

Fig. 10 depicts data structures used in support of a "popular sideyard" employed for popular-object storage;

Figs. 11A and 11B are diagrams that depict conventional train-algorithm object relocation;

Figs. 12A and 12B are diagrams similar to Figs. 11A and 11B but instead showing re-linking rather than reclamation of popular-object cars;

Fig. 13 is a diagram that illustrates the use of a hash table to store a remembered set used in support of the train algorithm;

Fig. 14 is a flow chart of a procedure for adding a new remembered-set entry;

Fig. 15 is a block diagram of a remembered-set structure;

5 Fig. 16 is a hash-table diagram similar to that of Fig. 13 but depicting contents that identify regions with coarser granularity;

Figs. 17A-D are diagrams that illustrate scanning the garbage-collected heap for references during a collection cycle that employs a multiple-car collection set; and

Fig. 18 is a flow chart of the procedure employed to search for references in a
10 region that a remembered-set entry identifies.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

The present invention eliminates much of the popular-object overhead by placing popular objects in their own cars. To understand how this can be done, consider
15 Fig. 9's exemplary data structures, which represent the type of information a collector may maintain in support of the train algorithm. To emphasize the ordered nature of the trains, Fig. 9 depicts such a structure 94 as including pointers 95 and 96 to the previous and next trains, although train order could obviously be maintained without such a mechanism. Cars are ordered within trains, too, and it may be a convenient to assign
20 numbers for this purpose explicitly and keep the next number to be assigned in the train-associated structure, as field 97 suggests. In any event, some way of associating cars with trains is necessary, and the drawing represents this by fields 98 and 99 that point to structures containing data for the train's first and last cars.

One such structure 100 is depicted as including pointers 101, 102, and 103 to
25 structures that contain information concerning the train to which the car belongs, the previous car in the train, and the next car in the train. Further pointers 104 and 105 point to the locations in the heap at which the associated car section begins and ends, whereas pointer 106 points to the place at which the next object can be added to the car section.

30 As will be explained in more detail presently, there is a standard car-section size used for all cars that contain more than one object, and that size is great enough to

contain a relatively large number of average-sized objects. But some objects can be too big for the standard size, so a car section may consist of more than one of the standard-size memory sections. Structure 100 therefore includes a field 107 that indicates how many standard-size memory sections there are in the car section that the structure manages—if the structure manages only a single car.

In the illustrated embodiment, that structure may instead manage a memory section that contains more than one (special-size) car section. When an organization of this type is used, one way of implementing the present invention involves including in structures like structure 100 a field 108 that indicates whether the heap space associated with the structure is used (1) normally, as a car section that can contain multiple objects, or (2) specially, as a region in which objects are stored one to a car in a manner that will now be explained by reference to the additional structures that Fig. 10 illustrates.

To deal specially with popular objects, the garbage collector may keep track of the number of references there are to each object in the generation being collected. Now, the memory space 110 allocated to an object typically begins with a header 112 that contains various housekeeping information, such as an identifier of the class to which the object belongs. To keep track of an object's popularity, the header can include a reference-count field 114. When the garbage collector locates references into the collection set during the collection cycle by processing the collection-set cars' remembered sets, it increments the reference-count field in the header of each collection-set object to which it finds a reference. Each time it does so, it tests the resultant value to determine whether the count exceeds a predetermined popular-object threshold. If so, it removes the object to a "popular sideyard" if it has not done so already.

Specifically, the collector consults a table 116, which points to linked lists of normal car-section-sized regions intended to contain popular objects. Preferably, the normal car-section size is considerably larger than the 30 to 60 bytes that has been shown by studies to be an average object size in typical programs. Under such circumstances, it would therefore be a significant waste of space to allocate a whole normal-sized car section to an individual object. For reasons that will become apparent below, collectors that follow the teachings of the present invention place each popular object

into its own, single car section. So the normal-car-section-sized regions to which table 116 points are to be treated as specially divided into car sections whose sizes are more appropriate to individual-object storage.

To this end, table 116 includes a list of pointers to linked lists of structures associated with respective regions of that type. Each list is associated with a different
5 object-size range. For example, consider the linked list pointed to by table 116's section pointer 118. Pointer 118 is associated with a linked list of normal-car-sized regions organized into *n*-card car sections. Structure 117 is associated with one such region and includes fields 120 and 122 that point to the previous and next structure in a
10 linked list of such structures associated with respective regions of *n*-card car sections. Car-section region 119, with which structure 117 is associated, is divided into *n*-card car sections such as section 124, which contains object 110.

More specifically, the garbage collector determines the size of the newly popular object by, for instance, consulting the class structure to which one of its header en-
15 tries points. It then determines the smallest popular-car-section size that can contain the object. Having thus identified the appropriate size, it follows table 116's pointer associated with that size to the list of structures associated with regions so divided. It follows the list to the first structure associated with a region that has constituent car sections left.

Let us suppose that the first such structure is structure 117. In that case, the
20 collector finds the next free car section by following pointer 126 to a car data structure 128. This data structure is similar to Fig. 9's structure 100, but in the illustrated embodiment it is located in the garbage-collected heap, at the end of the car section with which it is associated. In a structure-128 field similar to structure 100's field 129,
25 the collector places the next car number of the train to which the object is to be assigned, and it places the train's number in a field corresponding to structure 100's field 101. The collector also stores the object at the start of the popular-object car section in which structure 128 is located. In short, the collector is adding a new car to the object's train, but the associated car section is a smaller-than-usual car section, sized to
30 contain the newly popular object efficiently.

- 25 -

The aspect of the illustrated embodiment's data-structure organization that Figs. 9 and 10 depict provides for special-size car sections without detracting from rapid identification of the normal-sized car to which a given object belongs. Conventionally, all car sections have been the same size, because doing so facilitates rapid car identification. Typically, for example, the most-significant bits of the difference between the generation's base address and an object's address are used as an offset into a car-metadata table, which contains pointers to car structures associated with the (necessarily uniform-size) memory sections associated with those most-significant bits. Figs. 9 and 10's organization permits this general approach to be used while providing at the same time for special-sized car sections. The car-metadata table can be used as before to contain pointers to structures associated with memory sections whose uniform size is dictated by the number of address bits used as an index into that table.

In the illustrated embodiment, though, the structures pointed to by the metadata-table pointers contain fields exemplified by fields 108 of Fig. 9's structure 100 and Fig. 10's structure 117. These fields indicate whether the structure manages only a single car section, as structure 100 does. If so, the structure thereby found is the car structure for that object. Otherwise, the collector infers from the object's address and the structure's section_size field 134 the location of the car structure, such as structure 128, that manages the object's special-size car section, and it reads the object's number from that structure. This inference is readily drawn if every such car structure is positioned at the same offset from one of its respective car section's boundaries. In the illustrated example, for instance, every such car section's car structure is placed at the end of the car section, so its train and car-number fields are known to be located at predetermined offsets from the end of the car section.

To appreciate the effect that allocating popular objects to individual cars can have, consider the process of evacuating Fig. 11A's object A from car 1.1 during the collection of that car. (For consistency with Figs. 8A-J, the object symbols are again placed inside the car symbols, but, of course, the actual object data reside in the car sections that the car structures specify.)

Fig. 11A depicts object A as being referred to by object F, which resides in a car section associated with card 2.1, i.e. in a train newer than that in which object A re-

-26-

sides. The train algorithm therefore requires that object A be evacuated to object F's train. This could be done by moving object A's data into the car section associated with car 2.1. Another alternative, depicted in Fig. 11B, is to begin a new car in object F's train and place object A's data in the associated car section. This is the approach
5 conventionally taken when a previous car did not have enough room for the object being copied.

In either case, a result is that object F's reference to object A needs to be updated; a comparison of Figs. 11A and B reveals that object A is necessarily moved to a different memory-space location. But the actual physical copying—and thus the
10 pointer-update—are necessary in the illustrated case only because the car section 130 associated with object A's previous car 1.1 contains other objects, which will not in general end up in the car that object A newly occupies. In other words, if object A had started the collection cycle in its own car and were to end it in its own car, the only operation needed to place object A into a new car would be to relink the car data struc-
15 tures; no reference updates would be necessary. In the illustrated scenario, of course, updating object F's single pointer to object A would typically take less time than re-linking. For popular objects, though, the number of pointers requiring updating can be extremely large, so the benefit of changing car membership solely by relinking is significant.

20 The present invention therefore adopts the policy of placing popular objects in their own cars, as was mentioned above. Fig. 12A illustrates the result of this policy. For the sake of illustration, Fig. 12A shows region 130 as divided into smaller car sections sized to accommodate objects the size of object A. Fig. 12B illustrates moving object A to a new car, but it also indicates that object A remains stored in the same
25 memory-space location. Since object A has not moved, object F's reference to it remains the same, as do all other (not shown) references to object A.

Now, object A's reference count may be less than the popular-object threshold even though it is located in a popular-object car; i.e., it may have been placed in its popular-object car during a previous collection cycle, when its reference count was
30 higher. An advantageous feature of the present invention is that it permits such previ-

ously popular objects to be returned to regular cars. And a way of taking advantage of this capability involves selective re-linking of popular cars.

To appreciate the role that this re-linking plays in the collection cycle, it helps to review the overall sequence of a collection cycle. At the beginning of a collection
5 cycle, the card table is updated by scanning all cards marked by a write barrier as modified since the last collection cycle. For a generation that uses the train algorithm, this re-scanning of modified cards not only locates references into younger generations but also updates the appropriate remembered sets for cars in the same generation. As was mentioned above, one of the difficulties associated with popular objects is that they
10 tend to cause overly large remembered sets. In the illustrated embodiment, though, any remembered set associated with a popular-object car is allocated only a fixed amount of memory, and no further reference-identifying entries are made after the allocated memory is exhausted. Instead, the collector only updates a field that identifies the youngest train that contains a reference to the object in that car. Although the resultant remembered
15 set is incomplete, the collector operates in such a way that the object's memory space is not reclaimed prematurely. This will be described after a further review of the collection-cycle sequence.

After the card table is processed and the remembered sets updated, any younger generations are collected. For the sake of example, let us assume that the garbage-
20 collected heap is organized into two generations, namely, a young generation, which is completely collected during every collection cycle, and a mature generation, which is collected incrementally in accordance with the train algorithm. The young generation's collection may involve promotion of some objects into the mature generation, and a new train will typically be allocated for this purpose, unless an empty train already ex-
25 ists. When the young generation's collection is completed, collection of the generation organized in accordance with the train algorithm will begin, and at this point a new train will typically be allocated, too, unless the newest existing train is already empty.

The number of cars to be collected during this cycle is then determined in accordance with whatever policy the particular implementation uses. That number of cars
30 lowest in the sequence are identified as the collection set, possibly by removing them

from the trains to which they belong, linking them together in a single list, and setting per-car flags that indicate that they are parts of the collection set.

It is at this point that the re-linking referred to above occurs. The collector identifies the collection-set cars that are of the special popular-object variety. For each car thus identified, it checks the number of entries in the associated remembered set. As was mentioned above, the illustrated embodiment allocates only a fixed amount of memory space to each popular-object car's remembered set; if such a remembered set becomes too full, the collector stops adding entries to it. But it does keep updating a youngest-car value, associated with that car, that tells which is the youngest train that has a reference to that car's object. So, if the number of entries in the remembered set of a collection-set car identified as being of the popular-object variety is large enough to indicate that the collector may have omitted entries for further references, the collector assumes that the object contained in the car is still popular, and the car is immediately re-linked into the train that the youngest-car value identifies. Since the remembered set is to contain the addresses only of references from younger trains or younger cars in the same train, the collector accordingly empties the remembered-set list. Potentially, therefore, the associated remembered set will not be as full when the object comes up for collection again.

Note that, although the collector has assumed that the object is still popular, it does not really know that it is. The object's reference count is zero at the beginning of the collection cycle and not incremented until the updated remembered sets are processed, as has not yet happened. The collector therefore cannot rely on that value. Also, some of the remembered-set entries, on which the still-popular assumption was based, may be "stale"; this cannot be known, either, until the remembered-set entries are processed. On the other hand, an object may still be popular even if most of the remembered-set entries prove to be stale; there may be many other, "live" references for which the memory allocated to the remembered set had no further room for entries.

As was just explained, though, the collector nonetheless treats the object as though it is still popular and re-links its (popular-object) car section into a train young enough that the car section's remembered set can be cleared. (Remember, the illustrated embodiment's remembered set contains entries only for references from younger

trains.) If in fact the object had only a few surviving references, its remembered set will be small and no longer qualify it for the popularity assumption when that object next becomes part of a (subsequent) collection cycle's collection set—at least if it has not again become popular in the interim.

5 Although popular-object cars for which the number of remembered-set entries exceeds a threshold are linked into new trains immediately, as was just explained, the remainder stay in the collection set for processing with normal cars. That is, their remembered sets are processed to build per-train scratch-pad lists of references into the collection set. In the process, each collection-set object's reference count, which starts
10 at zero at the beginning of the collection cycle, is incremented each time a reference to it has been found. As was explained above, when all of the collection set's remembered sets have thus been processed (and cleared), the collector begins with the youngest train's scratch-pad lists of references and evacuates from the collection set into the respective train the objects to which the references refer.

15 If the object is in a normal car section but its reference count exceeds a predetermined threshold, then it is placed into a popular-object car section, as was explained above. On the other hand, if the reference count of an erstwhile popular object reveals that it is no longer popular, it is evacuated into a regular car section if it is not reclaimed. In order to introduce some hysteresis into popular-object car-section as-
20 signment, the threshold for indicating that an object is no longer popular will typically be made somewhat lower than the threshold used to determine that an object is newly popular.

 After the remembered-set processing, the collector processes any references from outside the generation into the collection set, as was also explained above, and all
25 collection-set objects thereby identified as having been referred to externally are placed into the youngest train.

 The scratch-pad list for the oldest train is then processed, after which the collector determines whether there are any references into the oldest train from trains outside it. If so, the collector reclaims any initial sequence of the current train's car sections into which no references remain. Otherwise, all of its car sections are reclaimed.
30 Preferably, the test for such references is performed not only for the oldest train but

also for other trains as well. One approach is to test the trains in order from the oldest to the youngest, reclaiming each train's car sections until a train is encountered into which references from outside it remain. This completes the collection cycle.

Popular objects are not the only ones whose placement into small, one-object
5 car sections is advantageous. In virtual-machine contexts, for instance, the interface from the virtual machine to so-called native methods not written in the virtual machine's language must frequently make copies of objects that are subject to relocation by the collector. But the interface can avoid this copying, which can be quite time-consuming in the case of large objects, if it is known that the collector will not relocate
10 the object. An example of a type of object for which this problem can arise frequently is the I/O buffer. Although such buffers are large, normal-size car sections tend to be much larger, so it would be wasteful to dedicate a whole normal-sized car section to an I/O buffer in order to avoid relocation and copying. Also, evacuation by copying takes more operations for larger objects than for smaller ones. Placing larger objects in one-
15 object sections reduces the cost of evacuating them, since it enables the evacuation to be done simply by re-linking, without re-locating.

In the illustrated embodiment, though, such waste is unnecessary, because "oversized" objects can be placed in car sections that are considerably smaller than those used for multiple objects and thereby remain stationary without wasting space
20 excessively. That is, an object can be a relatively small fraction of the normal car-section size in the illustrated embodiment and still be considered oversized. Such an object is stored in a special-size car section, just as a popular object is. However, it is marked as being oversized by, say, placing an appropriate entry in its car-structure field corresponding to Fig. 9's field 132. For instance, that field may be provided as a four-
25 byte field, with one of the bytes indicating whether the object is oversized and others indicating whether it is popular and whether it has only recently become so. (There are reasons not relevant to the present discussion why it may be convenient to distinguish newly popular objects from other popular objects.)

An oversized object thus placed in a special-sized car section is handled largely
30 in the manner described above for popular objects. Preferably, its remembered set, too, is allocated only a fixed amount of storage, so its car section is subjected to immediate

relinking when its remembered set gets too full, just as a popular object's is. The main difference is that an oversized object's reference count is not tested to determine whether the object will remain in the special car section; an object that starts out oversized remains oversized.

5 While the assignment of popular objects to their own cars eliminates the reference-update problem that popular objects present, it does nothing about another popular-object problem, which is the burden that such objects impose on the task of maintaining remembered sets. For popular-object cars, this burden is somewhat contained by the fact that remembered sets are not allowed to grow beyond a predetermined size.
10 But objects considerably less popular than the ones that qualify for their own cars may also be afflicted with this problem.

 To appreciate the remembered-set-maintenance problem, consider Fig. 13, which illustrates one of the many possible types of reference-list. There is theoretically little limit to a remembered set's size; if an object is referred to by every other object in
15 the generation, for instance, the remembered set for that object's car could be a significant fraction of the entire generation size. But it does not make sense to allocate that much space to each remembered set initially, and Fig. 13 depicts a memory space 150 allocated to the remembered set's reference list as containing only sixteen reference-sized locations.

20 At the beginning of each collection cycle, the collector inspects modified cards, as was mentioned above, and makes sure that any references from cars to cars farther forward in the queue are reflected in those farther-forward cars' remembered sets. That is, when the collector encounters such a reference, it should place the address of that reference into the remembered set of the referred-to object's car. But it should do so
25 only if the remembered set does not already contain that address: the collector should avoid duplicates.

 There are many approaches to achieving this result, and Fig. 13 illustrates one of them for the sake of concreteness. Let us suppose that the reference of interest occurs at a location whose address is 192E. To determine where to place this address in
30 the memory space 150 allocated to the reference list, the collector applies a hash function 152 to the address. In the illustrated example, the hash function is simply the ad-

- 32 -

dress's four least-significant bits, whose hexadecimal representation is E_H . The collector uses this value as an offset into the list, but it does not immediately store the address at the list location thus identified. It first reads that location's value to determine whether another address has already been stored there. In the Fig. 13 scenario, one already has.

Now, if that already-stored address were itself 192E, resulting from an entry made during a previous collection cycle, the collector would recognize that a duplicate had occurred, and it would stop its attempt to store the value. But the already stored address in the illustrated example is instead 477E, so the collector proceeds to the next reference-list location. This location, too, contains a non-NULL value that differs from the address to be stored. Since that location is at the end of the list, the collector proceeds circularly to the beginning of the list and repeats the test there. Again, the location is already occupied, so it proceeds still further, and this time it finds an empty location.

Even though the collector has not inspected every list entry, it can infer from encountering the empty location that the list has no duplicates of the entry to be added. Any previous attempt to store the same value would have taken the same steps, and the collector would accordingly have encountered the duplicate in attempting to find a space for this address. The collector therefore has simultaneously found a location and avoided duplication.

With the small list that Fig. 13 depicts, this approach to adding reference-list entries is fairly economical. But it does not scale particularly well. A car that includes a more-popular object may require a remembered-set reference list that is large indeed, and the amount of time taken to find a free location can significantly degrade performance. The illustrated embodiment provides two different ways of reducing this performance impact. In the illustrated embodiment, the previously described approach of employing fixed-size remembered sets is used for popular-object cars, while an approach about to be described is used for regular cars, although there is no reason in principle why the two approaches' uses need to be divided up in this manner.

The approach used for regular cars is to adjust remembered-set granularity dynamically. Fig. 14 is a flow chart that illustrates one way of entering new remembered-

set addresses in accordance with this approach. Block 154 represents beginning the process, and block 156 represents a test, which will now be described, for determining whether the space allocated to the list is already too full.

Now, a review of Fig. 13 reveals that the time required to find space for a new entry and avoid duplication increases disproportionately as the list becomes nearly full. To avoid this, the remembered-data structure may take a form such as the one that Fig. 15 illustrates. Whenever the collector makes an entry into the list, it increments the value of that structure's COUNT field 158 so that the field's contents represent the number of entries in the list. Fields 160 and 162 represent a threshold and the list's current capacity. The threshold, which may be, say, 70% of the capacity, is used to determine whether the list is too full for efficient operation.

If the count has not reached the threshold, the new entry is simply added, as Fig. 14's block 163 indicates. Otherwise, the collector increases the memory space allocated to the list, as block 164 indicates, unless the list size has already reached a limit value for which the step of block 166 tests. The size-increase operation of block 164 includes allocating more space to the list and re-entering the address values. For instance, the list size may be doubled and the hash-function output given an additional bit so that the address entries will be spaced more sparsely. This tends to reduce the time required to find an empty list location.

If the list size reaches the limit against which step 166 tests, the collector attempts to reduce the number of entries by increasing the granularity with which the entries specify reference locations. For example, the size of the region that each address entry is interpreted as specifying may change from only large enough to contain a single reference to one large enough to contain sixteen of them. This will be indicated by a change in a granularity indicator in Fig. 15's remembered-set field 168. To determine where references to the numbered set's car are located, the collector will consult the granularity entry and inspect the entire sixteen-location region for a reference into the car.

This granularity-coarsening operation, which block 170 represents, comprises masking the list's address entries and then entering the result. Fig. 16 illustrates this operation. Again, the address to be added is 192E, but in this case the granularity is

- 34 -

sixteen, so the last four bits ($2^4=16$) are masked to yield an entry value of 1920_H. The hash function 152 is changed to the least-significant four bits of the unmasked address portion, so the hash value is now 2_H.

The size of the table into which the hash value points is the same as that of the
5 Fig. 13 table, but the entries have been masked and stored in accordance with their respective new hash values. This usually results in fewer total entries. Let us assume, for instance, that the collector assigns a threshold value of twelve to a list whose capacity is sixteen. In Fig. 13, this limit has been exceeded: the number of entries is thirteen. With the granularity increase, the resultant number of entries depicted in Fig. 16 is
10 nine, which is less than the threshold value.

Of course, the time required to store any individual entry is not always less than before the granularity coarsening; as Fig. 16 shows, the collector actually has to inspect four entries before it finds that a further entry is unnecessary. But the delay still tends to be less on average.

Moreover, granularity coarsening can increase the effectiveness of optimizations such as single-entry caches. As Fig. 15 shows, the illustrated embodiment's remembered-set data structure includes a LAST_ENTRY field 174. This serves as a single-entry cache, into which the collector places a copy of the address entry that it has last placed in the list. Before the collector attempts to add an entry to the list, it first
20 compares the LAST_ENTRY value with the masked value of the address to be added. It may thereby determine that the address to be added is a duplicate.

For instance, note that the fourteenth location in Fig. 13 has the contents 192D. Suppose that this was the last entry made before the current one. After masking, that value is 1920, which is the same as the masked value of the new entry. So comparison
25 with the single-element cache 174 immediately reveals that the new entry would be a duplicate, and the collector therefore does not need to process it further.

Fig. 16 depicts the situation in which coarsening the granularity has opened up enough space in the list to add entries efficiently. But a review of the transition between Fig. 13 and Fig. 16 also reveals that this is not a necessary result. Fig. 14's
30 block 175 accordingly represents the step of determining whether enough space has in

fact resulted from the coarsening operation. If not, one could coarsen the granularity still further, but the embodiment that Fig. 14 illustrates instead restores the previous granularity, as block 176 indicates, and then increases the size, in a step represented by block 164. This results in a reasonable compromise between storage requirements and
5 the time required to locate references in a region that a remembered-set address-list entry specifies.

Although the above description of the Fig. 13 situation implies that each remembered set begins with single-byte granularity, it will probably be found preferable in most applications to start with a granularity somewhat coarser than this. This en-
10 ables a technique that will now be described to be employed to its best advantage.

The description of an example train-algorithm scenario set forth above in connection with Figs. 8A-J made the simplifying assumption that each collection cycle collected only a single car. But it is preferable for the collection-set size to be dynamically adjustable in accordance with current conditions. Strategies for doing this differ,
15 but one example is to set collection-set size to a value that approximately equals the average amount of space recently allocated between collection cycles. The fixed normal car size used in support of such a strategy would equal the increment by which collection-set size can be adjusted.

Now, the reason for maintaining a separate remembered set for each car is that
20 it cannot be known ahead of time which cars will be grouped together in a given collection set. Once the collection-set size is determined for a given collection cycle, though, the above-described process of evacuating objects from the collection set can largely be performed as though all of the collection-set objects occupy the same car.

But the fact that they do not actually occupy the same car tends to impose inefficiencies. In particular, since the different cars' remembered sets were necessarily
25 maintained separately, they are quite likely to specify reference-containing regions redundantly. That is, a region scanned once in response to a remembered-set entry for one of the collection-set cars may also be specified by another collection-set car's remembered set. Clearly, re-scanning the same region is wasteful. But attempts to avoid
30 this waste are complicated by the fact that the different remembered-set entries may

well specify reference-containing regions with different granularities. The illustrated embodiment provides a way of dealing with this complication.

Specifically, an array of Boolean values is associated with respective segments of the generation, and the collector marks the array during collection as it scans the corresponding regions to which the collection-set remembered-set entries direct it. Once a
5 segment is marked, the collector does not again scan it. Figs. 17A-D illustrate this approach. In Fig. 17A, we assume that the collection set has three cars, with which three remembered sets 180, 182, and 184 are respectively associated. As was mentioned above, all entries in the same remembered set identify their associated regions with the
10 same granularity, but we will also assume here that the three remembered sets' granularities are all different. Specifically, the farthest forward car's remembered set 180 has a medium granularity, the remembered set 182 of the car next in line has a relatively coarse granularity, and the last car's remembered set 184 has a relatively fine granularity.

For reasons that will become apparent as the description proceeds, it is preferable for the collector to begin collection with the address entries in the remembered set having the coarsest granularity, i.e., in remembered set 182. Fig. 17A depicts one of the remembered set 182's address entries ADDR[k] together with the region in the generation's memory space 60 to which that address points. The region is relatively large
20 because of the coarse granularity that remembered set 182 has acquired.

As was mentioned above, the collector proceeds largely as though all collection-set objects are in the same car. That is, when it inspects the region to which remembered set 182's address-list entry points, it searches not only for references into the car section associated with that remembered set car but also for references into the car sections managed by the cars associated with the other remembered sets 180 and 184.
25 When it finds them, it proceeds in the manner described above for the case of the single-car collection set.

Additionally, the collector keeps track of where it has scanned. It does so by making marks in a Boolean array 186 that contains an entry for each of the memory-space segments into which Fig. 17A shows the generation's memory space 60 as being
30 divided. As Fig. 17A illustrates, all of array 186's elements begin the collection cycle

-37-

with the same value, depicted in Fig. 17A as zero. Once the collector has finished inspecting the region to which remembered set 182's illustrated address-list entry ADDR[k] points, the collector enters a second value, which Fig. 17B depicts as one, in the array elements that correspond to all of the memory-space segments within the region referred to by remembered set 182's address-list entry. It is preferable but not
5 necessary for this segment size to equal that of the previously mentioned cards.

When the collector has processed all of remembered set 182's entries, it proceeds to remembered set 180, since that set has the next-coarsest granularity. In processing each of the address-list entries, it first consults the array to determine which
10 segments in the associated region have already been searched. By processing the remembered sets in the order of their granularities, the illustrated embodiment can expedite this process if the regions identified by the remembered-set entries are aligned in accordance with their granularity, that is, if the address of each region's first byte is the same, modulo the region's size. Specifically, these conditions enable the collector to
15 determine the values of all those segments' array elements by inspecting only a single array element—say, the first one—that represents a memory-space segment in the region to which the address-list entry refers.

In the case of entry ADDR[i], that entry has been switched to the value that indicates that the region has already been inspected for references into the collection set.
20 The collector thereby knows not to re-inspect the indicated region, because it has already done so in connection with remembered set 182. On the other hand, the entry in array 186 that corresponds to the first segment in the region to which ADDR[j] refers contains a value indicating that it has not yet been processed, so the collector proceeds to inspect that region for references into the collection set. When it does so, it records
25 the fact by entries in array 186, as Fig. 17C indicates.

Finally, the collector proceeds to process the finest-granularity remembered set 184. In this case, the granularity is finer than the segmentation on which the array 186 is based. By properly masking remembered set 184's address-list entries, though, the collector can still identify the encompassing segment. It thereby finds that
30 the region to which ADDR[m] points has already been inspected, so it does not re-inspect that region.

But a similar review of the array element corresponding to the segment encompassing the region to which ADDR[n] points finds that the segment has not been inspected. The collector accordingly proceeds to inspect that segment for references into the collection set. Since the region that the collector thereby inspects does not span the entire segment to which it belongs, though, the collector refrains from marking that segment as inspected, as Fig. 17D indicates. If there were a further fine-granularity remembered set that has an address-list entry identical to ADDR[n], the collector would not be apprised of the repetition, and it would accordingly re-examine the indicated region. So most implementations of this feature will be so designed that their minimum granularities will correspond to the segment size used by their arrays for indicating which regions have already been scanned.

The illustrated embodiment expedites collection-set processing even for collection cycles in which the collection set consists of only one car. According to this feature, the collector can abbreviate or eliminate the inspections it must perform in response to many of the collection-set address-list entries. This feature, too, is based on dividing the generation into fixed-size segments. In particular, when the collector is searching a remembered-set-identified region for references into the collection set, it does not have to inspect any of that region's segments that have been identified, in a manner presently to be described, as having no references into cars as far forward as the collection set's.

The collector maintains for each such segment at least an oldest-train value, preferably as part of an farthest-forward-car value. That farthest-forward-car value identifies the farthest-forward (same-generation) car to which a reference in that segment refers. The segments used for this purpose need not be of the same size as cards depicted in Fig. 6 for tracking inter-generation references. Indeed, this feature can be used by collectors that do not use the card-table approach to inter-generation-reference tracking. But we will use the term *cards* to refer to the segments with which the collector associates the farthest-forward-car values, because it is convenient to use the same size for both purposes.

In particular, the collector updates its farthest-forward-car values at the beginning of each collection cycle for all of a generation's cards that a write barrier em-

ployed for that purpose has identified as having been modified since the last collection cycle. So the write barrier used to identify cards whose modifications may necessitate card-table or remembered-set updates can also be used to identify those that may need their farthest-forward-car values updated. And the card-inspection operation used to re-
5 summarize those cards' inter-generation references can additionally include steps for updating the farthest-forward-car values. In such implementations, in other words, the farthest-forward-car value can be thought of as another field in each of Fig. 6's card-table entries.

In one exemplary approach, the farthest-forward-car value identifies the oldest
10 train and farthest-forward car by taking the form of a pointer to the farthest-forward car section's car structure. (In embodiments in which the "farthest-forward-car" value is actually used only to identify the oldest train, it may take the form of a pointer to the train structure.) When a card is first allocated, this pointer's value is NULL. It remains NULL so long as scans of that card triggered by that card's modifications between col-
15 lection cycles detect no references to cars farther forward than the one that includes that card.

When the write barrier has marked that card as modified since the last collection cycle, the collector sets that card's farthest-forward-car value to NULL before scanning it. Then, when it finds a reference into a farther-forward car, it not only updates that
20 referred-into car section's remembered set, as was mentioned above, but also updates the scanned card's farthest-forward-car value.

Specifically, the collector compares the train and car values indicated by the referred-to car section's car structure with the train and car values indicated by the car structure to which the scanned card's current farthest-forward-car value points. If that
25 comparison indicates that the referred-into car section is farther forward than the car section identified by the current farthest-forward-car value—or if the current farthest-forward-car value is NULL—the collector replaces the scanned card's current farthest-forward-car value with a pointer to the referred-into car section's car structure.

When the collector is to search a remembered-set-identified region for refer-
30 ences into the collection set, it follows a procedure like the one that Fig. 18 depicts. As block 190 indicates, the collector begins with the first card in the region that the re-

membered-set entry identifies. Before it inspects the card, it consults the farthest-forward-car value that it has maintained for this card. If every collection-set car is farther forward than the car thereby identified, it can skip the current card and proceed to any further cards in the region, as blocks 192, 194, and 196 indicate. Blocks 192
5 and 198 indicate that the card needs to be searched only when a card's farthest-forward-car value identifies a car at least as far forward as a car in the collection set. Actually, this is a slight oversimplification, as will be explained after a description of a particularly advantageous approach to making block 192's determination.

This approach involves having placed collection-set indicators in collection-set
10 car sections' car structures. When a collector implementing this approach has identified collection-set members at the beginning of a collection cycle, it makes the various data-structure changes necessary to remove the collection-set car sections from their trains, and it provides a collection-set indicator in each collection-set car structure by setting a Boolean `is_in_collection_set` field (not shown in the drawings) that car struc-
15 tures employed by implementations of this approach include. Then, if the farthest-forward-car value involved in the block 192 determination points to a car structure containing such an indicator, that determination is negative, and the collector scans the associated card. If that car structure's `is_in_collection_set` field is not set, on the other hand, the determination is positive, and the associated card is not scanned.

20 The description above of the block 192 determination is slightly oversimplified, because the farthest-forward-car value may not point to any car at all. It may be NULL, for instance, indicating that the card contains no references into a farther-forward car. In that case, of course, the collector does not scan the associated card. But there may also be other reserved values that some embodiments recognize as not
25 being valid pointers to car structures. For example, there may be a value, which we will call UNKNOWN, that the collector interprets as indicating that the card may contain references into cars farther forward but that the farthest-forward-car value does not point to the associated car structure. In that case, the collector would scan the associated card. In cases in which the region specified by the remembered-set entry contains
30 the whole card, the collector may use that card scan additionally to update the farthest-forward car entry.

To see how an UNKNOWN value may occur, consider the situation at the end of a collection cycle. Any car section that was in the collection set has just been either reclaimed or relinked, so it is unlikely to be the farthest forward car section referred into from within a card whose farthest-forward-car value still points its car structure.

5 That is, any such farthest-forward-car value would likely be invalid. But to update such farthest-forward-car values at the end of the collection cycle would be wasteful, since their associated cards may just end up being modified before the next collection cycle and require updating again.

A preferred approach is therefore simply to set such cards' farthest-forward-car
10 values to UNKNOWN. Now, some embodiments may update any UNKNOWN farthest-forward-car value at the beginning of the next cycle, just as though its card had been modified. But a card that has references to objects whose cars have just been collected is more likely than others not to have references into farther-forward cars, so there is some economy in omitting such a step. If that approach is taken, the collector
15 will typically update the value only at the point, if any, at which its card is actually encountered during remembered-set processing.

The frequency with which the UNKNOWN value is encountered can be reduced by providing for more than one farthest-forward car value for each card. When the "top" one or more of such values become invalid, the remaining values would
20 "move up," and the lowest values replaced with UNKNOWN. The UNKNOWN values would be replaced during updating, possibly with NULL if there are not enough referred-into cars to fill the farthest-forward car-value stack.

Some embodiments may expedite the block 192 determination further by immediately setting to UNKNOWN the farthest-forward-car value associated with a card
25 whose farthest-forward-car value points to a car structure that is determined in the block 192 step to be associated with a car section in the collection set. Although it is preferable to use the above-described use of the Boolean array to avoid repetitive scanning of the same region in a given collection cycle, not all collectors will use that feature. And it may be necessary to visit the same card more than once in a cycle even for
30 embodiments that do. If a remembered-set entry specifies a region smaller than a card, for instance, the Boolean-entry approach will not necessarily prevent the region-

containing card from being revisited. Without more, the collector would therefore have to repeat the process of following the farthest-forward-car value to the car structure and checking that structure for the collection-set indicator before it could conclude that it does indeed need to scan the region that the remembered-set entry identified.

5 By immediately setting the farthest-forward-car value to UNKNOWN, though, the collector avoids the need to repeat inspection of the car structure if that card is encountered again, since the UNKNOWN value will cause it to search in the associated card without doing so. Collectors that implement this feature will ordinarily leave the farthest-forward-car value equal to UNKNOWN in those situations in which the resultant scan reveals that the farthest forward car is indeed in the collection set. Otherwise,
10 they will update the farthest-forward-car value by making it equal to a pointer to the farthest forward car.

Note that the car farthest forward is not necessarily the one that is oldest chronologically. A comparison of Figs 8A and 8B illustrates this distinction. Those
15 drawings show that car 3.2 came into existence later than car 4.2. But car 3.2 is farther forward than car 4.2, because it belongs to an older train. Indeed, the general approach can be implemented by using only oldest-train values, without also maintaining a value representing that oldest train's farthest-forward car containing a reference to an object in the associated segment.

20 One of this feature's advantages is that it tends to avoid searching in response to "stale" remembered-set entries. Recall in this connection that the approach to remembered-set maintenance described above normally updates a remembered set at the beginning of a collection cycle only when inspection of a card marked by the write barrier as modified reveals that the card has a reference into the car with which the remembered set is associated. But inspecting a modified card identifies only cars into which it
25 now refers, not cars into which it previously referred. So this type of remembered-set update can only add entries; it cannot remove them. Remembered-set entries can therefore become stale. But use of farthest-forward-car values as just described reduces the number of searches through regions identified by stale remembered-set entries.

30 The illustrated embodiment also takes advantage of this fact in another way. The collector culls a remembered set before that set's associated car is collected. It

does so by scanning the farthest-forward-car values associated with the cards that the remembered set's entries identify. If no car specified by the farthest-forward-car values associated with any of the cards in a given region thus identified is at least as far forward as the remembered set's car, the collector discards the entry that identified that
5 region.

Now, one may initially conclude that this use of the farthest-forward-car values for advance culling is largely superfluous. After all, their use during the collection process virtually eliminates the stale entry's adverse effects. But further reflection reveals that this culling operation has a separate benefit. For one thing, even collectors that do
10 not use the farthest-forward-car test during collection can use it for advance culling.

More important, though, is that such culling can improve the remembered-set update process. To appreciate this, consider the remembered-set-entry process that Fig. 14 illustrates. Its blocks 164 and 170 respectively represent increasing the remembered set's size and coarsening its granularity. Each of those steps includes copying
15 entries from the space previously allocated to the set, possibly revising the entries, and placing the possibly revised entries into the remembered set's newly allocated space. So each such process presents a convenient opportunity to use the relevant farthest-forward-car values to ensure that no entry written into the new space is stale. In many cases, such culling will enable the collector to avoid or delay a further size increase or
20 granularity coarsening.

By making the sizes used for some single-object car sections smaller than those used for car sections that contain multiple objects, the present invention allows popular-object handling to be simplified greatly without significantly increasing the potential for perpetuating unreachable objects. It therefore constitutes a significant advance in
25 the art.

What is claimed is:

CLAIMS

- 1 1. In the method of garbage collection that treats a generation of a collected heap
2 as divided into car sections that belong to trains, so links the car sections as to order
3 them within the trains to which they belong, and collects the generation in collection
4 cycles, in each of which a collection set of at least one car section is collected in accor-
5 dance with the train algorithm, the improvement comprising:
 - 6 A) for each object in the collection set that is not in a popular-object car
7 section, determining whether the object meets a first popular-object cri-
8 terion;
 - 9 B) if any object is thereby determined to meet the first popular-object crite-
10 rion, designating at least one car section a popular-object car section;
 - 11 C) placing into a popular-object car section any object thereby determined
12 to meet the first popular-object criterion;
 - 13 D) storing no more than one object in a popular-object car section; and
 - 14 E) when a popular-object car section containing a reachable object meeting
15 the popular-object criterion is collected, collecting it by re-linking the
16 car section without relocating the object.
- 1 2. A method as defined in claim 1 wherein:
 - 2 A) the popular-object car sections are of different predetermined car-section
3 sizes; and
 - 4 B) for a popular object that will fit in a popular-object car section, the
5 placement of an object in a popular-object car section comprises placing
6 the object in a car section whose size is the smallest of the car-section
7 sizes that can contain the object.
- 1 3. A method as defined in claim 2 wherein:
 - 2 A) the method further comprises maintaining a plurality of popular-object-
3 car-section lists of popular-object-car-section regions, each of which is
4 associated with a different one of the car-section sizes and contains
5 popular-object car sections of the associated car-section size; and

6 B) the placement of an object in a popular-object car section comprises
7 placing the object in the first popular-object-car-section region in the
8 popular-object-car-section list that is associated with the smallest car-
9 section size that can contain the object.

1 4. A method as defined in claim 3 wherein:

2 A) the method further includes storing a plurality of objects in each of at
3 least one car section of a standard-car-section size; and
4 B) the size of the popular-object-car-section regions equals the standard-
5 car-section size.

1 5. A method as defined in claim 2 further including:

2 A) determining whether an object contained in a popular-object car section
3 meets a second popular-object criterion; and
4 B) reclaiming or relocating to a car section that is not a popular-object car
5 section any object thereby determined not to meet the second popular-
6 object criterion.

1 6. A method as defined in claim 5 that includes determining whether an object
2 contained in a popular-object car section meets a second popular-object criterion only if
3 the object fails to meet a third popular-object criterion.

1 7. A method as defined in claim 5 wherein:

2 A) the method further includes counting references to objects; and
3 B) the first popular-object criterion is that the number of references counted
4 for an object exceeds a first reference-count value.

1 8. A method as defined in claim 7 wherein the second popular-object criterion is
2 that the number of references counted for an object exceed a second reference-count
3 value.

1 9. A method as defined in claim 8 that includes determining whether an object
2 contained in a popular-object car section meets a second popular-object criterion only if
3 the object fails to meet a third popular-object criterion.

1 10. A method as defined in claim 9 wherein the third popular-object criterion is that
2 the remembered set associated with a popular-object car section that contains the object
3 contains more than a predetermined number of remembered-set entries.

1 11. A method as defined in claim 7 wherein the first and second reference-count
2 values differ.

1 12. A method as defined in claim 5 wherein the first and second popular-object cri-
2 teria differ.

1 13. A method as defined in claim 1 further including:
2 A) determining whether an object contained in a popular-object car section
3 meets a second popular-object criterion; and
4 B) reclaiming or relocating to a car section that is not a popular-object car
5 section any object thereby determined not to meet the second popular-
6 object criterion.

1 14. A method as defined in claim 13 that includes determining whether an object
2 contained in a popular-object car section meets a second popular-object criterion only if
3 the object fails to meet a third popular-object criterion.

1 15. A method as defined in claim 13 wherein:
2 A) the method further includes counting references to objects; and
3 B) the first popular-object criterion is that the number of references counted
4 for an object exceeds a first reference-count value.

1 16. A method as defined in claim 15 wherein the second popular-object criterion is
2 that the number of references counted for an object exceed a second reference-count
3 value.

1 17. A method as defined in claim 16 that includes determining whether an object
2 contained in a popular-object car section meets a second popular-object criterion only if
3 the object fails to meet a third popular-object criterion.

1 18. A method as defined in claim 17 wherein the third popular-object criterion is
2 that the remembered set associated with a popular-object car section that contains the
3 object contains more than a predetermined number of remembered-set entries.

1 19. A method as defined in claim 15 wherein the first and second reference-count
2 values differ.

1 20. A method as defined in claim 13 wherein the first and second popular-object
2 criteria differ.

1 21. A garbage collector that treats a generation of a collected heap as divided into
2 car sections that belong to trains, so links the car sections as to order them within the
3 trains to which they belong, and collects the generation in collection cycles, in each of
4 which a collection set of at least one car section is collected in accordance with the train
5 algorithm, wherein the garbage collector:

6 A) for each object in the collection set that is not in a popular-object car
7 section, determines whether the object meets a first popular-object crite-
8 rion;

9 B) if any object is thereby determined to meet the first popular-object crite-
10 rion, designates at least one car section a popular-object car section;

11 C) places into a popular-object car section any object thereby determined to
12 meet the first popular-object criterion;

13 D) stores no more than one object in a popular-object car section; and

14 E) when a popular-object car section containing a reachable object meeting
15 the popular-object criterion is collected, collects it by re-linking the car
16 section without relocating the object.

1 22. A garbage collector as defined in claim 21 wherein:

2 A) the popular-object car sections are of different predetermined car-section
3 sizes; and

4 B) for a popular object that will fit in a popular-object car section, the
5 placement of an object in a popular-object car section comprises placing
6 the object in a car section whose size is the smallest of the car-section
7 sizes that can contain the object.

1 23. A garbage collector as defined in claim 22 wherein:

2 A) the garbage collector further maintains a plurality of popular-object-car-
3 section lists of popular-object-car-section regions, each of which is asso-
4 ciated with a different one of the car-section sizes and contains popular-
5 object car sections of the associated car-section size; and

6 B) the placement of an object in a popular-object car section comprises
7 placing the object in the first popular-object-car-section region in the
8 popular-object-car-section list that is associated with the smallest car-
9 section size that can contain the object.

1 24. A garbage collector as defined in claim 23 wherein:

2 A) the garbage collector further stores a plurality of objects in each of at
3 least one car section of a standard-car-section size; and

4 B) the size of the popular-object-car-section regions equals the standard-
5 car-section size.

1 25. A garbage collector as defined in claim 22 that further:

2 A) determines whether an object contained in a popular-object car section
3 meets a second popular-object criterion; and

4 B) reclaims or relocates to a car section that is not a popular-object car sec-
5 tion any object thereby determined not to meet the second popular-
6 object criterion.

1 26. A garbage collector as defined in claim 25 that determines whether an object
2 contained in a popular-object car section meets a second popular-object criterion only if
3 the object fails to meet a third popular-object criterion.

1 27. A garbage collector as defined in claim 25 wherein:

2 C) the garbage collector further counts references to objects; and

3 D) the first popular-object criterion is that the number of references counted
4 for an object exceeds a first reference-count value.

1 28. A garbage collector as defined in claim 27 wherein the second popular-object
2 criterion is that the number of references counted for an object exceed a second refer-
3 ence-count value.

1 29. A garbage collector as defined in claim 28 that determines whether an object
2 contained in a popular-object car section meets a second popular-object criterion only if
3 the object fails to meet a third popular-object criterion.

1 30. A garbage collector as defined in claim 29 wherein the third popular-object cri-
2 terion is that the remembered set associated with a popular-object car section that con-
3 tains the object contains more than a predetermined number of remembered-set entries.

1 31. A garbage collector as defined in claim 27 wherein the first and second refer-
2 ence-count values differ.

1 32. A garbage collector as defined in claim 25 wherein the first and second popular-
2 object criteria differ.

1 33. A garbage collector as defined in claim 21 that further:

2 A) determines whether an object contained in a popular-object car section
3 meets a second popular-object criterion; and

4 B) reclaims or relocates to a car section that is not a popular-object car sec-
5 tion any object thereby determined not to meet the second popular-
6 object criterion.

1 34. A garbage collector as defined in claim 33 that determines whether an object
2 contained in a popular-object car section meets a second popular-object criterion only if
3 the object fails to meet a third popular-object criterion.

1 35. A garbage collector as defined in claim 33 wherein:

2 A) the garbage collector further counts references to objects; and

3 B) the first popular-object criterion is that the number of references counted
4 for an object exceeds a first reference-count value.

1 36. A garbage collector as defined in claim 35 wherein the second popular-object
2 criterion is that the number of references counted for an object exceed a second refer-
3 ence-count value.

1 37. A garbage collector as defined in claim 36 that determines whether an object
2 contained in a popular-object car section meets a second popular-object criterion only if
3 the object fails to meet a third popular-object criterion.

1 38. A garbage collector as defined in claim 37 wherein the third popular-object cri-
2 terion is that the remembered set associated with a popular-object car section that con-
3 tains the object contains more than a predetermined number of remembered-set entries.

1 39. A garbage collector as defined in claim 35 wherein the first and second refer-
2 ence-count values differ.

1 40. A garbage collector as defined in claim 33 wherein the first and second popular-
2 object criteria differ.

1 41. A storage medium containing instructions readable by a computer to configure
2 the computer to operate as a garbage collector that treats a generation of a collected
3 heap as divided into car sections that belong to trains, so links the car sections as to or-
4 der them within the trains to which they belong, and collects the generation in collec-
5 tion cycles, in each of which a collection set of at least one car section is collected in
6 accordance with the train algorithm, wherein the garbage collector:

- 7 A) for each object in the collection set that is not in a popular-object car
8 section, determines whether the object meets a first popular-object crite-
9 rion;
- 10 B) if any object is thereby determined to meet the first popular-object crite-
11 rion, designates at least one car section a popular-object car section;
- 12 C) places into a popular-object car section any object thereby determined to
13 meet the first popular-object criterion;
- 14 D) stores no more than one object in a popular-object car section; and
- 15 E) when a popular-object car section containing a reachable object meeting
16 the popular-object criterion is collected, collects it by re-linking the car
17 section without relocating the object.

1 42. A storage medium as defined in claim 41 wherein:

- 2 C) the popular-object car sections are of different predetermined car-section
3 sizes; and
- 4 D) for a popular object that will fit in a popular-object car section, the
5 placement of an object in a popular-object car section comprises placing
6 the object in a car section whose size is the smallest of the car-section
7 sizes that can contain the object.

1 43. A storage medium as defined in claim 42 wherein:

- 2 C) the garbage collector further maintains a plurality of popular-object-car-
3 section lists of popular-object-car-section regions, each of which is asso-

4 ciated with a different one of the car-section sizes and contains popular-
5 object car sections of the associated car-section size; and

6 D) the placement of an object in a popular-object car section comprises
7 placing the object in the first popular-object-car-section region in the
8 popular-object-car-section list that is associated with the smallest car-
9 section size that can contain the object.

1 44. A storage medium as defined in claim 43 wherein:

2 C) the garbage collector further stores a plurality of objects in each of at
3 least one car section of a standard-car-section size; and

4 D) the size of the popular-object-car-section regions equals the standard-
5 car-section size.

1 45. A storage medium as defined in claim 42 wherein the garbage collector further:

2 C) determines whether an object contained in a popular-object car section
3 meets a second popular-object criterion; and

4 D) reclaims or relocates to a car section that is not a popular-object car sec-
5 tion any object thereby determined not to meet the second popular-
6 object criterion.

1 46. A storage medium as defined in claim 45 wherein the garbage collector deter-
2 mines whether an object contained in a popular-object car section meets a second
3 popular-object criterion only if the object fails to meet a third popular-object criterion.

1 47. A storage medium as defined in claim 45 wherein:

2 E) the garbage collector further counts references to objects; and

3 F) the first popular-object criterion is that the number of references counted
4 for an object exceeds a first reference-count value.

1 48. A storage medium as defined in claim 47 wherein the second popular-object
2 criterion is that the number of references counted for an object exceed a second refer-
3 ence-count value.

1 49. A storage medium as defined in claim 48 wherein the garbage collector deter-
2 mines whether an object contained in a popular-object car section meets a second
3 popular-object criterion only if the object fails to meet a third popular-object criterion.

1 50. A storage medium as defined in claim 49 wherein the third popular-object crite-
2 rion is that the remembered set associated with a popular-object car section that con-
3 tains the object contains more than a predetermined number of remembered-set entries.

1 51. A storage medium as defined in claim 47 wherein the first and second refer-
2 ence-count values differ.

1 52. A storage medium as defined in claim 45 wherein the first and second popular-
2 object criteria differ.

1 53. A storage medium as defined in claim 41 wherein the garbage collector further:
2 A) determines whether an object contained in a popular-object car section
3 meets a second popular-object criterion; and
4 B) reclaims or relocates to a car section that is not a popular-object car sec-
5 tion any object thereby determined not to meet the second popular-
6 object criterion.

1 54. A storage medium as defined in claim 53 wherein the garbage collector deter-
2 mines whether an object contained in a popular-object car section meets a second
3 popular-object criterion only if the object fails to meet a third popular-object criterion.

1 55. A storage medium as defined in claim 53 wherein:
2 C) the garbage collector further counts references to objects; and
3 D) the first popular-object criterion is that the number of references counted
4 for an object exceeds a first reference-count value.

1 56. A storage medium as defined in claim 55 wherein the second popular-object
2 criterion is that the number of references counted for an object exceed a second refer-
3 ence-count value.

1 57. A storage medium as defined in claim 56 wherein the garbage collector deter-
2 mines whether an object contained in a popular-object car section meets a second
3 popular-object criterion only if the object fails to meet a third popular-object criterion.

1 58. A storage medium as defined in claim 57 wherein the third popular-object crite-
2 rion is that the remembered set associated with a popular-object car section that con-
3 tains the object contains more than a predetermined number of remembered-set entries.

1 59. A storage medium as defined in claim 55 wherein the first and second refer-
2 ence-count values differ.

1 60. A storage medium as defined in claim 53 wherein the first and second popular-
2 object criteria differ.

1 61. A computer data signal embodied in a carrier wave and representing sequences
2 of instructions that, when executed by a processor, cause it to operate as a garbage col-
3 lector that treats a generation of a collected heap as divided into car sections that belong
4 to trains, so links the car sections as to order them within the trains to which they be-
5 long, and collects the generation in collection cycles, in each of which a collection set
6 of at least one car section is collected in accordance with the train algorithm, wherein
7 the garbage collector:

- 8 A) for each object in the collection set that is not in a popular-object car
9 section, determines whether the object meets a first popular-object crite-
10 rion;
- 11 B) if any object is thereby determined to meet the first popular-object crite-
12 rion, designates at least one car section a popular-object car section;
- 13 C) places into a popular-object car section any object thereby determined to
14 meet the first popular-object criterion;

- 15 D) stores no more than one object in a popular-object car section; and
16 E) when a popular-object car section containing a reachable object meeting
17 the popular-object criterion is collected, collects it by re-linking the car
18 section without relocating the object.

1 62. A computer data signal as defined in claim 61 wherein:

- 2 A) the popular-object car sections are of different predetermined car-section
3 sizes; and
4 B) for a popular object that will fit in a popular-object car section, the
5 placement of an object in a popular-object car section comprises placing
6 the object in a car section whose size is the smallest of the car-section
7 sizes that can contain the object.

1 63. A computer data signal as defined in claim 62 wherein:

- 2 A) the garbage collector further maintains a plurality of popular-object-car-
3 section lists of popular-object-car-section regions, each of which is asso-
4 ciated with a different one of the car-section sizes and contains popular-
5 object car sections of the associated car-section size; and
6 B) the placement of an object in a popular-object car section comprises
7 placing the object in the first popular-object-car-section region in the
8 popular-object-car-section list that is associated with the smallest car-
9 section size that can contain the object.

1 64. A computer data signal as defined in claim 63 wherein:

- 2 A) the garbage collector further stores a plurality of objects in each of at
3 least one car section of a standard-car-section size; and
4 B) the size of the popular-object-car-section regions equals the standard-
5 car-section size.

- 56 -

1 65. A computer data signal as defined in claim 62 wherein the garbage collector
2 further:

- 3 A) determines whether an object contained in a popular-object car section
4 meets a second popular-object criterion; and
5 B) reclaims or relocates to a car section that is not a popular-object car sec-
6 tion any object thereby determined not to meet the second popular-
7 object criterion.

1 66. A computer data signal as defined in claim 65 wherein the garbage collector
2 determines whether an object contained in a popular-object car section meets a second
3 popular-object criterion only if the object fails to meet a third popular-object criterion.

1 67. A computer data signal as defined in claim 65 wherein:

- 2 A) the garbage collector further counts references to objects; and
3 B) the first popular-object criterion is that the number of references counted
4 for an object exceeds a first reference-count value.

1 68. A computer data signal as defined in claim 67 wherein the second popular-
2 object criterion is that the number of references counted for an object exceed a second
3 reference-count value.

1 69. A computer data signal as defined in claim 68 wherein the garbage collector
2 determines whether an object contained in a popular-object car section meets a second
3 popular-object criterion only if the object fails to meet a third popular-object criterion.

1 70. A computer data signal as defined in claim 69 wherein the third popular-object
2 criterion is that the remembered set associated with a popular-object car section that
3 contains the object contains more than a predetermined number of remembered-set en-
4 tries.

1 71. A computer data signal as defined in claim 67 wherein the first and second ref-
2 erence-count values differ.

1 72. A computer data signal as defined in claim 65 wherein the first and second
2 popular-object criteria differ.

1 73. A computer data signal as defined in claim 61 wherein the garbage collector
2 further:

3 A) determines whether an object contained in a popular-object car section
4 meets a second popular-object criterion; and

5 B) reclaims or relocates to a car section that is not a popular-object car sec-
6 tion any object thereby determined not to meet the second popular-
7 object criterion.

1 74. A computer data signal as defined in claim 73 wherein the garbage collector
2 determines whether an object contained in a popular-object car section meets a second
3 popular-object criterion only if the object fails to meet a third popular-object criterion.

1 75. A computer data signal as defined in claim 73 wherein:

2 A) the garbage collector further counts references to objects; and

3 B) the first popular-object criterion is that the number of references counted
4 for an object exceeds a first reference-count value.

1 76. A computer data signal as defined in claim 75 wherein the second popular-
2 object criterion is that the number of references counted for an object exceed a second
3 reference-count value.

1 77. A computer data signal as defined in claim 76 wherein the garbage collector
2 determines whether an object contained in a popular-object car section meets a second
3 popular-object criterion only if the object fails to meet a third popular-object criterion.

1 78. A computer data signal as defined in claim 77 wherein the third popular-object
2 criterion is that the remembered set associated with a popular-object car section that

3 contains the object contains more than a predetermined number of remembered-set en-
4 tries.

1 79. A computer data signal as defined in claim 75 wherein the first and second ref-
2 erence-count values differ.

1 80. A computer data signal as defined in claim 73 wherein the first and second
2 popular-object criteria differ.

1 81. A garbage collector that treats a generation of a collected heap as divided into
2 car sections that belong to trains, so links the car sections as to order them within the
3 trains to which they belong, collects the generation in collection cycles, in each of
4 which a collection set of at least one car section is collected in accordance with the train
5 algorithm, and comprises:

6 A) means for determining, for each object in the collection set that is not in
7 a popular-object car section, whether the object meets a first popular-
8 object criterion;

9 B) means for designating at least one car section a popular-object car sec-
10 tion if any object is thereby determined to meet the first popular-object
11 criterion;

12 C) means for placing into a popular-object car section any object thereby
13 determined to meet the first popular-object criterion;

14 D) means preventing any more than one object from being placed in a
15 popular-object car section; and

16 E) means for, when a popular-object car section containing a reachable ob-
17 ject meeting the popular-object criterion is collected, collecting it by re-
18 linking the car section without relocating the object.




FIG. 1

2/28




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6

6/28




FIG. 7




FIG. 8A




FIG. 8B

9/28




FIG. 8C




FIG. 8D




FIG. 8E







FIG. 8F


92

3.2	B, L
4.1	O
4.2	
5.1	

92


4.1	O
4.2	
5.1	
5.2	





92

4.2	4.3	5.1	5.2
	O		L






FIG. 9




FIG. 10




FIG. 11A




FIG. 11B




FIG. 12A




FIG. 12B




FIG. 13




FIG. 16

22/28




FIG. 14

23/28




FIG. 15

24/28




FIG. 17A

25/28




FIG. 17B

26/28




FIG. 17C

27/28




FIG. 17D

28/28




FIG. 18

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 00/22685

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F12/02 G06F9/44

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, PAJ, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	SELIGMAN ET AL: "Incremental Mature Garbage Collection Using the Train Algorithm" ECOOP. EUROPEAN CONFERENCE ON OBJECT-ORIENTED PROGRAMMING, XX, XX, vol. 952, 26 March 1995 (1995-03-26), pages 235-252, XP002105706 page 239, line 23 -page 240, line 17 ---	1-81
A	HUDSON ET AL: "Incremental Collection of Mature Objects" MEMORY MANAGEMENT. INTERNATIONAL WORKSHOP IWMM. PROCEEDINGS, XX, XX, September 1992 (1992-09), pages 388-406, XP002105707 page 401, line 1 - line 16 --- -/--	1,21,41, 61,81

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document but published on or after the international filing date

L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

* & * document member of the same patent family

Date of the actual completion of the international search

29 November 2000

Date of mailing of the international search report

06/12/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Nielsen, O

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 00/22685

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	HOSKING A L ET AL: "A COMPARATIVE PERFORMANCE EVALUATION OF WRITE BARRIER IMPLEMENTATIONS" ACM SIGPLAN NOTICES, ASSOCIATION FOR COMPUTING MACHINERY, NEW YORK, US, vol. 27, no. 10, 1 October 1992 (1992-10-01), pages 92-109, XP000327291 ISSN: 0362-1340 page 94, left-hand column, line 8 -page 96, right-hand column, line 19 ---	1,21,41, 61,81
A	WO 98 48353 A (SUN MICROSYSTEMS INC) 29 October 1998 (1998-10-29) page 4, line 28 -page 6, line 18 ---	1,21,41, 61,81
P,A	EP 0 940 755 A (SUN MICROSYSTEMS INC) 8 September 1999 (1999-09-08) page 7, paragraph 39 -page 9, paragraph 48 -----	1,21,41, 61,81

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 00/22685

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9848353 A	29-10-1998	US 5953736 A EP 0914634 A	14-09-1999 12-05-1999
EP 0940755 A	08-09-1999	CN 1236921 A JP 2000047931 A	01-12-1999 18-02-2000