

Rochester Institute of Technology

RIT Digital Institutional Repository

Theses

3-1-1983

Garbage Collection of Linked Data Structures: An Example in a Network Oriented Database Management System

Delphine T H Kung

Follow this and additional works at: <https://repository.rit.edu/theses>

Recommended Citation

Kung, Delphine T H, "Garbage Collection of Linked Data Structures: An Example in a Network Oriented Database Management System" (1983). Thesis. Rochester Institute of Technology. Accessed from

This Thesis is brought to you for free and open access by the RIT Libraries. For more information, please contact repository@rit.edu.

Rochester Institute of Technology
School of Computer Science and Technology

GARBAGE COLLECTION OF LINKED DATA STRUCTURES:
AN EXAMPLE IN A NETWORK ORIENTED DATABASE MANAGEMENT SYSTEM

by

DELPHINE T.H. KUNG

A thesis submitted to
The Faculty of the School of Computer Science and Technology
in partial fulfillment of the requirements for the
degree of Master of Science in Computer Science

MARCH, 1983

Approved by:

----- 5/11/84
Thesis Advisor Mrs. Mary Ann Dvorch Date

----- 5/17/84
Committee Member DOCTOR Peter H. Lutz Date

----- 5/17/84
Committee Member DOCTOR Rayno D. Niemi Date

G-929169

Statement for granting or denying permission to
reproduce an RIT thesis

Title of Thesis "Garbage collection of linked data structure:
an example in a network oriented database management system"

I Delphine Kung hereby deny permission to the Wallace Memorial
Library, of R.I.T., to reproduce my thesis in whole or in part.
Any reproduction will not be for commercial use or profit.

Signature: _____

May 3, 1984
date

ACKNOWLEDGEMENT

I would like to thank my thesis adviser, Mary Ann Dvorch, for her conduct guidance and discussion in directing my effort on this project. I would like to express my thanks to Dr. Rayno Niemi and Dr. Peter Lutz for the ideas they stimulated and the time they have generously spent during the evolution of this project. Lastly, I would like to place the most gratitude to my parents for their encouragement and support letting me concentrate on my study and research in accomplishing this work.

ABSTRACT

A unified view of the numerous existing algorithms for performing garbage collection of linked data structure has been presented. An implementation of a garbage collection tool in a network oriented database management system has been described

Table of Content

| | |
|--|----|
| 1. INTRODUCTION | 5 |
| 2. THEORETICAL AND CONCEPTUAL DEVELOPMENT | 8 |
| 2.1. Collecting Single-sized Cells | 14 |
| 2.1.1. Marking | 14 |
| 2.1.2. Reclaiming Marked cells | 21 |
| 2.1.3. Moving Collectors | 23 |
| 2.2. Collecting Varisized Cells | 26 |
| 2.2.1. Marking | 26 |
| 2.2.2. Reclaiming Marked Cells | 28 |
| 2.2.3. Moving Collectors | 33 |
| 3. SPECIAL TOPICS ON GARBAGE COLLECTION | 34 |
| 3.1. Collecting in Virtual Memory | 34 |
| 3.2. Reference Counter | 38 |
| 3.3. Parallel and Real-Time Collections | 43 |
| 3.4. Language Implementation | 49 |
| 3.5. Final Remarks | 54 |
| 4. AN IMPLEMENTATION OF GARBAGE COLLECTOR IN A NETWORK DBMS | 60 |
| 4.1. Problem Statement | 61 |
| 4.2. Implementation Part Approach | 62 |
| 4.2.1. Database Files | 63 |
| 4.2.1.1. Record Types | 65 |
| 4.2.1.2. Set Types | 66 |
| 4.2.2. Database Deletion | 68 |

| | |
|--------------------------------|----|
| 4.2.2.1. Delete Record | 69 |
| 4.2.2.2. Delete Member | 69 |
| 4.2.2.3. Delete Owner | 70 |
| 4.2.3. Prototype | 70 |
| 4.2.4. Algorithms | 78 |
| 5. BIBLIOGRAPHY | 80 |

1. INTRODUCTION

Computer programs and languages having a need for the dynamic creation of data objects must return the data objects' space that is no longer used to a pool of free storage. Hence the objects' unused storage may be reused in the future for the creation of new data objects. Garbage collection is the method that has been traditionally used to do this bookkeeping. The process of garbage collection is one of reclaiming unused storage space which can be accomplished by various algorithms. Experience with large LISP programs indicates that substantial execution time - 10 to 30 percent - is spent in garbage collection [STEELE 75, WADLER 76]. Hence the efficiency of programs depends directly on the availability of fast methods for garbage collection.

Garbage collection has become important. Every data structure book has devoted sections to discuss this topic [KNUTH 73, HOROWITZ & SAHNI 76, PFALTZ 77, GOTLIEB 78, AUGENSTEIN 79, STANDISH 80], Knuth's book, section 2.3.5, is the most comprehensive. It contains detailed descriptions and analyses of some of the garbage collection algorithms that appeared prior to 1968, and, despite its age, it remains a standard reference for algorithms. Although numerous papers have been published since the late sixties, very little has been written to summarize and

classify current work. The purpose of this thesis is to investigate a unified view of the many existing algorithms for performing garbage collection of linked data structures. This thesis is divided into two parts: research and implementation. The objectives of the research portions are:

- (1) to review the classical algorithms for collecting linked data structures, and to discuss their marking and collecting phases, with and without compacting;
- (2) to provide a unified description of recent garbage collection algorithms and to explain how they relate to the classical algorithms;
- (3) to survey the related topics of parallel and real-time garbage collections, the use of reference counters, the use of secondary and virtual storage, and language features which influence the design of collectors.
- (4) to present a comprehensive bibliography on the subject for the interested reader.

Although storage allocation and garbage collection are interrelated, the emphasis of this research is on garbage collection proper, that is, on reclaiming storage. Buddy systems [KNUTH 73] and related work are not covered.

The implementation portion of the thesis explores the use of a garbage collection tool in a network oriented Database Management System [DBMS] which uses three major data structures: B-trees, binary trees and doubly linked list.

2. THEORETICAL AND CONCEPTUAL DEVELOPMENT

A cell is a number of contiguous computer words that can be made available to a user. Cells are requested by the user's program from the storage allocator. Since the number of available cells is finite, a time may come when no single cell is available.

Experience indicates that some of the cells are inaccessible and unused but still occupy space in the memory. A cell becomes "garbage" when it can no longer be accessed through the pointers of any reachable cell. It is the garbage collector's task to reclaim this wasted storage space, and return it to the allocator.


Typically, a garbage collector is invoked automatically when the free storage list is near exhaustion. Users never know, unless they become suspicious because the computer seems to drop dead suddenly, but temporarily. This is a static garbage collection. At this time, processing is stopped and the system takes over, going through an exhaustive and often time-consuming process of determining exactly what list nodes are still in use, and which ones can be returned to free storage. All these free nodes are returned, and then user processing resumes with a replenished supply of free storage [see KNUTH 73]. However, some DFMS such as SEED [see SEED 81] are equipped

with a tool that has to be called to identify the percentage of uncovered freespace. The Data Base Administrator [DBA] uses this utility to monitor areas that are recognized as highly volatile to avoid last minute scrambles for space. The DBA then calls the garbage collector to reclaim freespace if the percentage of uncovered freespace is high enough to warrant garbage collection. (SEED is a DBMS based on the DBMS specification given in the April 1971 CODASYL DBTG report.)


Garbage collection is usually called automatically either when the allocator runs out of space or shortly before. Higher level languages often contain primitives for requesting groups of words from the allocator. The garbage collector may be turned on when one of these primitives is executed. For example, in LISP, some newer garbage collectors do not use the mark-and-sweep idea. Instead they do a little garbage collection with each "cons". They never stop for a complete collect. The function "cons" also calls the garbage collector [see WINSTON 81, pp.106-110].

The difficulty with the static garbage-collection technique is that it requires an execution time proportional to the number of nodes in use. It tends to be slow when the memory is nearly full. And in such cases the number of free storage cells found by the reclamation

process is not worth the effort. This unattractive performance under pressure is a major drawback of garbage collection [KNUTE 73]. A further problem is that an allocation may be made from the middle of a physically contiguous area of storage (from the end of the first of two contiguous blocks, the second of which is not yet known to be free). Thus unnecessary fragmentation can result. In diagram (1), the heavily shaded areas at the extreme left and right are unavailable. Suppose that there are 2 adjacent areas of memory, both of which are available, but because of the garbage-collection philosophy, one of them (shown shaded) is not in the available space list.



We may now reserve the areas known to be available (shown on diagram (2)):



If garbage collection occurs at this point, we have two separate free areas (shown on diagram (3)).

```

-----
// //|      |/////|      |////////
-----

```

(3)

And as time goes on the situation gets worse. This phenomenon causes the static garbage collection technique to leave memory more broken up than it should be. If we had used a philosophy of returning blocks to the available list as soon as they become free (real-time dynamic garbage collectio), and collapsing adjacent available areas together, we have collapsed diagram (1) into diagram (4).

```

-----
////////|      |////////
-----

```

(4)

and we would reserve the areas known to be availble on diagram (5).

```

-----
////////|      |//////////
-----

```

(5)

In the case discussed above, the situation portrayed by Diagram (5) is much better than the situation protrayed by Diagram (3).

To remove this difficulty of storage fragmentation, we use marking together with the process of compacting memory, i.e., moving all the reserved blocks into the consecutive location; so that all available blocks will come together whenever garbage collection is done. The allocation now becomes completely trivial, since there is only one available block. Even though this technique takes time to recopy all the location that are in use, and to change the value of the link fields therein, it can be applied efficiently with a disciplined use of pointers. However, garbage collection is unsuitable for real-time applications, because even if the garbage collector goes into action infrequently, it requires large amounts of computer time on these occasions.

After my investigation into the matter, I have chosen to categorize garbage collection into two separate phases as follows:

- (1) the identification of the storage space that may be reclaimed.
- (2) the collection of this reclaimed space and its return to the memory area available to the user.

Two principal methods have been suggested for phase one:

- (a) keep counters that indicate the number of times cells

have been referenced, and identify the inaccessible cell count as zero. This identification method is called the reference counter.

(b) keep a list of accessible cells and trace their links and mark every accessible cell. This identification method is usually called marking.

Phase 2 generally can be classified into 2 approaches:

(a) incorporate all available cells that are linked by pointers into a free list.

(b) compact all used cells in one end of memory. All freed cells are made available to the allocator in the other end.

There are various ways to do the compaction according to the relative positions of the cells left after compaction. It can be classified into the following three types:

(b-1) Arbitrary: Cells which originally point to one another do not necessarily occupy contiguous positions in the memory after compaction.

(b-2) Linearizing: Cells which originally point to one another become adjacent in the memory after compaction.

(b-3) Sliding: Cells are moved toward one end of the storage space without changing their linear order.

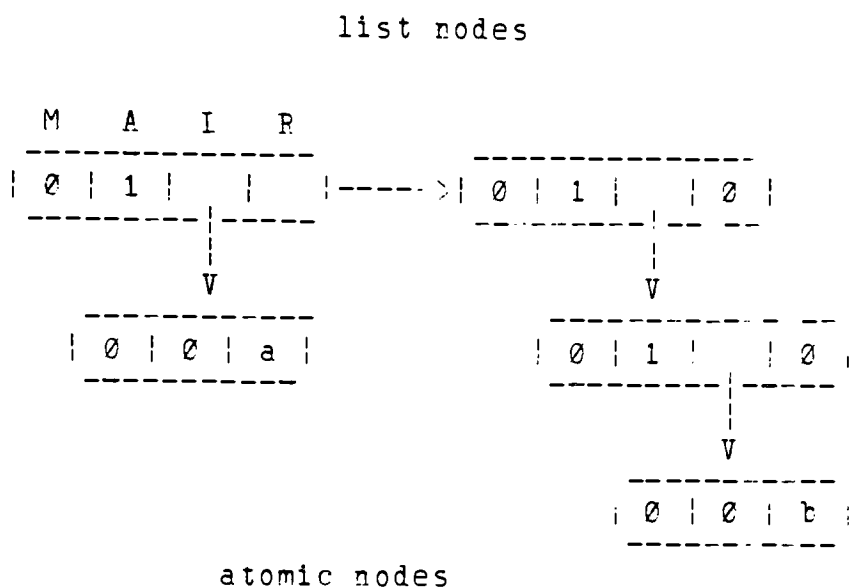
It is also convenient to classify garbage collection according to the type of cells that are reclaimed. The early methods were applicable only to programs in which all the cells were of the same sizes. With the introduction of records (or similar structures) into programming languages, it became important to do garbage collection in programs involving cells of different sizes.

2.1. Collecting Single-sized Cells

2.1.1. Marking

LISP cells illustrate the problems involved in marking single-size cells. Each LISP cell has two fields: the left (or car) and the right (or cdr). These fields contain pointers either to other cells or to atoms (special kinds of cells containing no pointers). Each cell also contains two Boolean fields: one to help differentiate between atomic and nonatomic cells, and the other to be used in marking. With this node structure, the list(a,(b)) is represented in Figure 2-1 [referenced from HOROWITZ 76]:

The algorithm shown in Figure 2-2 is a recursive procedure generally used for marking LISP lists (including



M: Mark bit M = 0 (unmarked) M = 1 (marked)
 A: Atom bit A = 0 (atomic node) A = 1 (list node)
 I: Left pointer
 R: Right pointer

Figure 2-1

atoms). It uses three auxiliary procedures:

nonatomic(p): Boolean function that tests whether the cell pointed to by p is nonatomic;

unmarked(p): Boolean function that tests whether the cell pointed to by p is unmarked;

marknode(p): Procedure that marks the cell pointed to by p that turned on its marking bit (marking bits are initially

turned off)).

Note the similarity of the marking algorithm in Figure 2-2 with the classical preorder tree-traversal algorithm [see KNUTH 73]. The one in Figure 2-2, however, can handle general lists, including circular ones.

The following is perhaps the simplest and oldest marking algorithm in use. A 1-bit field MARK (initially 0) for marking is associated with each node. A stack S (initially empty) is used by the algorithm to hold references to sublists in use but not yet marked. The algorithm marks an unmarked list by marking each element in turn, then marking each unmarked sublist of the list. This algorithm is quite simple and straight forward [see KNUTH 73]. A nonrecursive version of the algorithm in

```

procedure mark(p)    /* p is pointer that is called
                      by value */
begin
  if unmarked(p)
  then
    marknode(p)
    if nonatomic(p)
    then
      mark(left(p));
      mark(right(p))
end.

```

Figure 2-2

Figure 2-2 could be also designed with this stack technique. It has, however, one major drawback. As I have indicated, garbage collection is being performed because storage is essentially exhausted, and yet the collector requires, for its execution, storage for a stack of indeterminate size. For example, if the storage area consists of n IISP cells, the maximum depth required for the stack is then n . To reserve this much additional storage initially is uneconomical.

A fixed-size stack can of course be reserved for the routine, with the space taken away from that available for list storage. But if that stack is exhausted during garbage collection, there is no recourse. Partial reclamation does not work because nodes not yet processed are all apparently not in use - their MARK bits are 0. This strategy fails if the stack overflows. Several algorithms have been proposed to eliminate this difficulty; all of them involve reducing the required storage by trading it for longer time needed in which to do the marking. Many interesting algorithms have been defined to effect marking [referenced from KNUTH '73, pp. 414-419].

The first of these algorithms (see Algorithm C in KNUTH '73, p.415) uses a stack of fixed length h , where h is smaller than n . (Suppose that the storage area consists of n IISP cells.) The pointers to the cells being

marked are stacked using mod h as the stack index. In other words, the stack can be thought of as being "circular", and when its index exceeds h , the additional information is written over previously stored information. The stack therefore only "recalls" the most recently stored h items and "forgets" the other ones. First, the immediately accessible cells (i.e., the cells pointed to by certain fixed locations in the main program which are used as a source for all memory accesses.) are marked. Marking then proceeds as in the algorithm in Figure 2-2. However, since some cells that should have been "remembered" have been "forgotten", the stack will become empty before the task is complete. When this happens, the memory is scanned from the lowest address, looking for any marked cell whose contents point to unmarked cells. If such a cell is found, marking resumes as before and continues until the stack becomes empty again. Eventually, a scan will find no marked cells referring to unmarked cells, and marking is complete. Actually, the scanning need not start from the beginning of the memory each time. The next scan may begin just after the recorded last address of the previous scan.

An elegant alternative to the use of a separate stack is the temporary use of the link fields of the lists themselves, to contain the recovery information held in the

stack. These fields are modified during the marking process to allow the list traversals, then reset to their proper values as the marking of each list is completed. This strategy is known as back-tracing which is an elegant marking method developed independently by Peter Deutsch and by Herbert Schorr and W.M. Waite in 1965. As the algorithm moves to the right and down through sublists, the link fields are temporarily set in such a way that the algorithm can later retrace its steps back along the list. This garbage collection method is very generally applicable even though it is not fast. Its execution time is still proportional to the number of nodes marked. This technique requires one additional bit per cell. The additional bit per cell (called a tag bit) indicates the direction in which the restoration of reversed links should proceed (i.e., whether to follow the left or the right pointer). Knuth suggests a method for avoiding using a tag bit by instead using the bit (atom bit) already necessary for testing whether a cell is atomic. [HOROWITZ 76] contains clear descriptions of this method.

Veillon [see VEILLON 76] has shown that it is possible to transform the classical recursive algorithm in Figure 2-2 into the Deutsch/Schorr-Waite algorithm. First, the parameter of the recursive procedure "mark" is eliminated by introducing the link reversal feature. The two

recursive calls of the resulting parameterless procedure, needed to mark the left and right fields, are eliminated by introducing the tag bits to differentiate between the returns from the two calls.

Knuth proves by induction (from Algorithm A to Algorithm E) the correctness of the link-reversal marking algorithm of Deutsch/Schorr-Waite. An alternate proof may be obtained by noting that the transformations suggested by Veillon preserve correctness. Wegbreit [WEGBREIT 72] proposes a change of the Deutsch/Schorr-Waite algorithm that uses a bit stack instead of a tag bit per cell. In the light of Veillon's program transformation, one sees that Wegbreit's stack simply implements the returns from the parameterless recursive procedure derived from Figure 2-2.

Schorr and Waite then proposed using a hybrid algorithm which combines a fixed-size stack with their link-reversal technique. It consists of using the stack algorithm whenever possible. If stack overflow occurs, the tracing and marking proceed by the method of link reversal [see KNUTH 73, p. 592].

We have seen that, at most, three bits per cell are necessary to perform IISP's garbage collection. The first two are used in recognizing atoms and in marking; the

third one is used as a tag bit, if needed by the algorithm. It should be pointed out that these three bits need not be located within or near their corresponding cells. special areas of the memory (bit maps or tables) may be allocated for this purpose. Whether or not this should be done is machine dependent.

The algorithms described in this section can be generalized to cover cells of a single-size m , with $m > 2$. The generalized version of the algorithm in Figure 2-2 would involve recursively marking each of the m fields of the cell.

2.1.2. Reclaiming Marked Cells

The simplest method for reclaiming the marked cells consists of linearly sweeping the entire memory. After turning off their mark bit, unmarked cells are incorporated into the free list administered by the storage allocator.

If compacting is preferred, it can be performed by scanning the memory twice. In the first scan, two pointers are used, one starting at the bottom of the memory (higher address), the other at the top. The top one is incremented until it points to an unmarked cell;

the bottom pointer is then decremented until it points to a marked cell. The contents of the marked cell are thereupon moved to the unmarked cell. A pointer to the new cell is placed in the old cell so that the new cell can be found during the second scan. The mark bit is also turned off during this scan. By the time the two pointers meet, all marked cells have been compacted in the upper part (lower address) of the memory.

. The second scan is needed for readjusting the pointers since some cells have been moved. It is essential to update any pointers to obsolete cell locations. This scan sweeps only the compacted area. Since we placed the pointer to the new address in the obsolete cell. Pointers to the obsolete cell locations are readjusted whenever they point to cells whose contents have been moved from the liberated area to the compacted area of the memory. In other words, each of these pointers is replaced by the new pointer. According to Knuth, this method was first proposed by D. Edwards [KNUTH 73, pp. 421]. LISP and AIGOL 60 programs describing in detail this method of compacting have appeared in [COHEN 67]. Note that the two-pointer compactor is of the arbitrary type; after compaction, cells that originally point to one another do not necessarily occupy contiguous positions of the memory.

2.1.3. Moving Collectors

An obvious algorithm for garbage collection would be to output all useful (i.e., reachable) data to the secondary storage area and then to read them back to the main memory. This, however, has several drawbacks [see MINSKY 63]:

- (1) It may require additional storage equally as large as the main memory.
- (2) The time overhead for transferring between memories is usually considerable.
- (3) Unless special precautions are taken, shared cells would be output more than once, in which case the main memory may not be sufficiently large for reading back the information (this situation becomes critical when the main memory contains loops of pointers).

Minsky [MINSKY 63] proposes an algorithm that eliminates the difficulties described in (3). His algorithm does not use a stack, but requires one marking bit per IISP cell. Each unmarked cell is traced and marked. Triplets (the new address of a cell and the content of its left and right fields) are computed and output to the secondary storage. The new address is also placed in the marked cell, and whenever a pointer to that cell is

encountered, the pointer is adjusted to reflect the move. When the triplets are subsequently read back into the main memory, the contents of the fields are stored in the specified new address. Minsky's algorithm has the advantage of compacting the useful information into one area of the main memory. After compaction, list elements that are linked are positioned next to each other, making Minsky's algorithm a linearizing compactor. These two properties are very important when virtual memory is used, as will be discussed in Chapter 3.

In Minsky's algorithm, fields of the original list are used to store information about the output list which is in the secondary storage. Consequently, the original list is destroyed. In this respect, it is convenient to distinguish between the terms moving and copying. The former implies a possible destruction of the original structure, whereas the latter does not. Minsky's algorithm can be used to move lists in contexts other than garbage collection. Since its appearance, several other algorithms have been proposed to do moving or copying. They can be used for garbage collection purposes as well. Most are designed to move or to copy lists without resorting to mark bits or to a stack. As in Minsky's algorithm, (1) a forwarding address is usually left in the old cell, and pointers referring to that cell are readjusted accord-

ingly, and (2) the moved lists are compacted in contiguous positions of the memory.

A few algorithms have been proposed for coping lists without using a stack or mark bits. They differ from the moving algorithms in which the altered contents of old lists are later restored to their original values. Clark [CLARK 75, CLARK 78] and Fisher [FISHER 75] discuss this kind of coping of trees and general lists.

Fenichel and Yochelson [FENICHEL 69] suggest a variant of Minsky's collector which uses an implicit stack but does not require mark bits. They divide the available memory into two areas called semispaces. At a given time, only one area is used by the allocator. When its space is exhausted, the reachable lists are moved to the other space in a linearized compacted form. The algorithm is intended for use in a paging environment.

Cheney's algorithm [CHENEY 70, WAIDEN 72], Reingold's algorithm [REINGOLD 73], and Clark's algorithm [CLARK 76, GOTLIEB 78] all represent improvements over the previous algorithm: they require neither a stack nor mark bits. Cheney's algorithm is done by moving the list to a contiguous area. Reingold's algorithm is achieved by using the Deutsch/Schorr-Waite link-reversal technique mentioned in Section 2.1.1.. And Clark's algorithm moves a list

into a contiguous area of the memory with the stack implicit in the list being moved. Clark shows that his algorithm is in most cases more efficient than both Cheney's and Reingold's.

2.2. Collecting Varisized Cells

2.2.1. Marking

As I have indicated, garbage collection in programs involving varisized cells became important with the introduction of records (or similar structures) into programming languages. Figure 2-3 shows a marking algorithm similar to that of Figure 2-2, but applicable to varisized cells. Two additional auxiliary procedures are used:

number(p): an integer function yielding the number of contiguous words (items) in the cell to which p points (this information is stored in the cell itself); and

field(p,i): a function yielding the *i*th item of the cell pointed to by p.

It is assumed that p always points to the first item of the cell. The algorithm in Figure 2-3, like that in Figure 2-2, requires stack storage space when none may be available. If the memory contains *n* cells of various

```

procedure mark(p)      * p is a pointer that is called
                        by value */

integer i;

begin
  if unmarked(p)
    then
      marknode(p);
      if nonatomic(p)
        then
          for i=1 until number(p) do
            mark(field(p,i))
          end.

```

Figure 2-3

sizes, the maximum depth required for the stack is n . Two additional fields are reserved per cell for distributing stack storage among the cells. Essentially, these fields contain the quantities p and i needed to implement the recursive calls of the procedure in Figure 2-3. A description of an algorithm of this kind appears in [THORELLI 72].

The marking algorithms of Section 2.1.1. which use a fixed-length stack can also be adapted to process vari-sized cells. They may then use a fixed-length stack of height h with "stack index = mod h " as before, but each stack position will contain information corresponding to p and i in the algorithm of Figure 2-3.

Variants of the Deutsch/Schorr-Waite link-reversal algorithm applicable to varisized cells are described in [MARSHALL 71, WDCON 71, THORELLI 72, THORELLI 76 and HAN-SCN 77].

2.2.2. Reclaiming Marked Cells

The method of compacting described in Section 2.1.2 is not applicable to varisized cells, since marked and unmarked cells cannot be swapped if they are of different sizes.

A possible storage configuration made up of varisized cells is shown in Figure 2-4 [referenced from WAITE 73, pp. 178-183]. The "title" is used to reserve the first word of every cell, it contains the length of cell and its own address [detail in WAITE 73, pp. 178-183]. To move all cells to a compact block at the bottom of memory, several algorithms have been proposed. One of the earliest is that of Haddon and Waite [HADDON 67, WAITE 73]. This compactor is of the sliding type and performs two scans of the entire memory. The objective of the first scan is to perform the compaction and to build a "break table" [see Figure 2-5], which the second scan uses to readjust the pointers.




Figure 2-4: A possible configuration of a 2000-word store
(The hatched areas represent free storage)

This algorithm updates all pointers after the elements have been moved. As the elements are moved to a compact group at the lower end of the storage area, a table is built for use in relocation. For each set of consecutive blocks which are not free, this break table has one entry giving the address of the first word of the set and the number of free words which were found below the set in the uncompact store. At the end of the first

scan, the break table occupies the liberated part of the memory. It is then sorted on the address value to speed up the pointer readjustment done by the second scan. Readjustment consists of examining each pointer, consulting the table (using binary search) to determine the new position of the cell it used to point to, and changing the pointer accordingly [HADDON 67].

Figure 2-5 shows the break table corresponding to the storage of Figure 2-4. After all blocks have been moved, the relocatable address are updated by looking them up in the break table and subtracting the number of free words associated with that particular address. If an address is not equal to any table entry, the number of free words associated with the next-lower table entry is used. Any address lower than the lowest table entry is unchanged. Using Figure 2-5, for example, the relocatable address $POT+950$ would be altered to contain $POT+100$ ($= POT+950-850$) and $POT+990$ would become $POT+140$ ($= POT+990-850$). This procedure may require more processing time (depending upon the layout of the free blocks in memory). An interesting feature of Haddon and Waite's algorithm is that no additional storage is needed to build the break table since it can be proved that the space available in the free storage suffices to store the table.

| Address | Free words |
|----------|------------|
| BOT+100 | 100 |
| BOT+950 | 850 |
| BOT+1600 | 1250 |

Figure 2-5: The break table for Figure 2-4

Other compacting algorithms for varisized cells have been proposed. The LISP 2 garbage collection algorithm described in [KNUTH 73, p. 602-602], and those presented in [WEGPREIT 72] and [TECRELLI 76] have the following features in common.

Three linear scans are used. In the first scan the holes (inaccessible cells) are linked to form a free list. Two fields are reserved in each hole to store its size and a pointer to the next hole. (A subsequent scan may combine adjacent holes into a single larger hole.) The second scan consists of recognizing pointers and using the information contained in the free list to adjust them. Once the pointers have been readjusted, a third scan takes care of moving the accessible cells to the compacted area. This compactor is therefore of the sliding type.

The second scan, which readjusts pointers, is the most time consuming of the three scans. Wegbreit [WEGBREIT 72] proposes variants of this algorithm which make this scan more efficient. One variant consists of constructing a break table (called directory) which summarizes the information contained in the free list of holes. However, storage for the directory may be unavailable. Wegbreit suggests trying to use the largest hole for this purpose. When this is possible, binary search can speed up pointer readjustment. Many interesting variants for collecting varisized cells are described in [LANG 72, KNUTH 73, ZAVE 75, TERASHIMA 78, MORRIS 78 and MORRIS 79].

Marking, pointer readjustment, and compacting can be made simpler if the list processing "preserves address ordering". This means that nodes are allocated sequentially and placed into a contiguous section of memory. Under these conditions, marking can be performed in a single scan through the entire memory without using a stack. This scan also finds the number of active cells, which the second scan then uses for readjusting the pointers. The third and final scan performs the compaction which is of the sliding type. This method is applicable to fixed or variable size cells. Details are given in [FISHER 74].

2.2.3. Moving Collectors

Some of the moving algorithms mentioned in Section 2.1.3. may be adapted to handle varisized cells. A representative of this class of algorithms is described in the section 3.1 since it is particularly suitable for operation in virtual memory. An algorithm for copying varisized cells is described in [STANDISH 80]. It requires cells to have an additional field large enough to store an address. A first pass consists of linking all used cells via the additional field [see THOPFELI 72]. A second pass copies each cell c_i in the linkage and inserts the copy, c'_i , as the successor of c_i . The successor of c'_i becomes the cell c_{i+1} . After this copying, the odd-numbered elements of the new linkage contain the original cells and the even-numbered ones contain the copies. Finally, a third pass readjusts the pointers in each copied element and separates the copy from the original. Fisher [FISHER 75] and Robson [ROBSON 77] also describe algorithms for copying LISP cells that can be generalized for copying varisized cells.

3. SPECIAL TOPICS ON GARBAGE COLLECTION

3.1. Collecting in Virtual Memory

The ratio of the size of secondary memory to the size of main memory is an important factor in designing collectors that operate in virtual memory. When this ratio is small, some of the algorithms described in the previous sections may be used. The methods described in this section are suitable when the ratio is large.

The use of secondary storage through paging [COHEN 67] changes the design considerations for implementing garbage collection algorithms in important ways. First, it is no longer necessary to try to avoid using additional storage for a stack, since the size of the available virtual memory in current systems is considerable. Avoiding page faults and thrashing (caused by having structures whose cells are scattered in many pages) becomes a critical factor in improving the efficiency of garbage collection. Compaction is for this reason important when collecting in this environment. Cohen and Trilling show that garbage collection with compaction brings about significant time gains in performance of LISP programs. They also find that a direct transcription of the classical garbage collection algorithms to a virtual memory

environment can lead to unbearably slow collection times. [CIARK 79] contains additional useful information about the performance of compacting collectors operating in virtual memory. Compaction of cells in virtual memory should not only eliminate unused holes but should also construct the compacted area so that pointers refer, if possible, to neighboring cells. As mentioned in Section 2.1.3., Minsky's algorithm [see MINSKY 63] satisfies this requirement.

Measurements in actual IISP programs show that about 97 percent of list cells have just one reference to them [CIARK 77]. This property is important when designing garbage collection algorithms which operate in virtual memory.

Bobrow and Murphy [BOEROW 67] show that the use of a selective "cons" can improve the efficiency of subsequent processing and garbage collection. Basically, they advocate keeping one free-list per page. A new cell requested by a call of `cons[x,y]` is taken from the free area of a page according to the following strategy:

- (1) If possible, take from the page containing the cell pointed to by `y`; otherwise,
- (2) take from the page containing the cell pointed to by `x`; otherwise.

(3) take from the page containing the most recently created cell; otherwise,

(4) take from any page containing a fair number (say, 16 of free cells.

The purpose is to minimize page faults in manipulating linked lists. Additional information on garbage collection using virtual memory can be found in [BOERGA 67,68].

An important design consideration for implementing garbage collection algorithms in a paging environment is deciding when collection should be invoked. Since very large memories are currently available, it seems reasonable to collect whenever page faults render the program processing unbearably slow.

A class of algorithms suitable for use in virtual memory is the one described by Baker [BAKER 78]. It is based on the copying collector proposed by Benichou-Yochelson and by Cheney that was briefly described in Section 2.1.3.. What follows is a more detailed presentation of this type of algorithm. Although it is applicable in collecting varisized cells, this presentation applies only to IISP cells.

The available memory is divided into two areas called semispaces. At a given time, only one is used by the allocator. During garbage collection, the reachable lists are moved to the other space in a compacted form.

Bishop proposed an approach for designing collectors that operate in a very large virtual memory (of the order of 10^{12} bits). Even using the real-time approaches which will be discussed in Section 3.3 of this paper, it is impractical to garbage collect the entire memory at one stretch since large portions of memory may remain unchanged during program execution. Bishop suggests collecting only in parts of the address space rather than in the entire space. (A similar approach is used in Ross' AED system [ROSS 67].) The memory is divided into areas that can be collected independently, and a variant of the Fenichel and Yochelson collector is used. This collector increases the locality of reference, an important factor in a paging environment.

Tracing and copying are performed only within a given area. The system keeps lists of all area references, both incoming and outgoing. Incoming references are modified to point to the area's new copy; they define the immediately accessible cells from which collection starts. Before discarding the old copy of an area, its useless outgoing references are removed from the corresponding

memory. The disadvantages of this approach are:

- (1) the extra space needed for the counters,
- (2) the overhead required to update the counters, and
- (3) the inability to reclaim general cyclic structures.

The principal drawback of the reference counter method is the fact that it does not always free all the cells which are available. It cannot work for the circular list. For example, recursive lists will never be returned to storage by the reference counter technique. Their counts will be nonzero (since they refer to themselves) even when no other list accessible to the running program points to them [detail in HOROWITZ 76, pp. 165-169].

Deutsch, Knuth, and Weizenbaum suggest combining the reference counter technique with classical garbage collection in order to overcome the problems of each [see DEUTSCH 76, KNUTH 73, WEIZENBAUM 69]. The former would be used during most of the processing time; the latter, being more expensive, would be done as a last resort. This allows the use of small reference counter (thus reducing the storage requirements) because counters that reach their maximum value remain unmodified. Classical garbage collection, called when the free list is exhausted, starts

lists of incoming references.

Bishop developed a method for maintaining the lists of areas references and indicated that this can be done automatically without incurring substantial run-time overhead. He advocates altering the virtual memory mechanism to cause traps when interarea references are stored into cells, and shows how virtual memory hardware can be constructed to perform this extra service efficiently.

3.2. Reference Counter

The reference counter approach toward storage reclamation provides for constant monitoring of list linkages, with lists returned to free storage as soon as they become available. This is a dynamic garbage collection.

The use of reference counters [advocated by COLLINS 60 and WEIZENPAUM 63] has recently attracted renewed interest. An extra field, called refcount, is required for each cell to indicate the number of times the cell is referenced. This field has to be updated each time a pointer to the cell is created or destroyed. When the refcount becomes equal to zero, the cell is inactive and can be collected. At least, theoretically, the refcount must be large enough to hold the number of cells in the

of a program still active, and those which are truly unreferenced and can be reclaimed. It follows from (1) and (2) that if a cell's address is not in the MPT or the ZCT, its refcount is one.

(3) The variable reference table (VRT) contains the addresses of cells referred to by program variables (including the temporary variables in the recursion stack).

Deutsch and Bobrow note that a transaction file is built up consisting of three sorts of transactions, that may affect the accessibility of data. They are (1) allocation of a new cell from free space, (2) creation of a pointer to a cell, and (3) destruction of a pointer to a cell. Instead of updating the hash tables as the transactions occur, Deutsch and Bobrow propose storing them in a sequential file. The transactions are examined at suitable time intervals and then the tables are updated.

When a new cell is allocated, its address should be placed in the ZCT. Since this is usually followed by the creation of a pointer to the newly allocated cell (that implies removal from the ZCT), the pair of transactions can be ignored.

According to the hybrid approach, when a pointer is created, it is examined before its insertion into a cell

or pointer variable. Three cases are possible:

(a) The pointer refers to a cell in the MRT. The corresponding refcount value is then increased by one if it has not already reached its maximum; otherwise, the refcount value is left unchanged.

(b) The pointer refers to a cell in the ZCT. The cell is then removed from that table, since its count becomes one.

(c) If tests (a) and (b) fail, the pointer refers to a cell having a reference count of one. It must then be placed in the MRT with a refcount of two.

When a pointer is destroyed (removed from a cell), two cases are possible:

(a) the pointer refers to a cell in the MRT. The cell's refcount value is decreased by one, except when it has reached its maximum, in which case it is left unchanged. If the new value of reference count is one, the cell is removed from the MRT.

(b) The pointer does not refer to a cell in the MRT. Its count is one by default, and should be reduced to zero. The cell is therefore entered in the ZCT.

The VRT is used when incorporating new cells into the free list. Since the stack is constantly being updated,

the VRT is only computed periodically. A cell is reclaimed when its address is listed in the ZCT but not in the VRT. The ZCT is updated by eliminating the entries of the reclaimed cells that are not pointed to by the program variables.

For the classical collection, Deustch and Bobrow advocate using a variant of the two-semispace collector of Fenichel and Yockelson [FENICHEL 69]. They also point out that an auxiliary processor could speed up the collection. Its task would be to scan the ZCT and VRT tables to determine which cells could be incorporated into the free list.

Barth [BARTE 77] considers reference counters in relation to shifting garbage collection overhead to compile time. He shows that savings in collection time are sometimes possible by carefully studying, at compile time, the program's assignments. For example, if $r \leftarrow \text{right}(r)$, the cell originally pointed to by r may be incorporated into the free list if it is known that it will not be referenced by other pointers.

3.3. Parallel and Real-Time Collections

Two proposals have been made to circumvent the time lost by garbage collecting interruptions. The first is to

allow garbage collection to proceed simultaneously with program execution by using two parallel processors: one is responsible for collection, the other for program execution. When collection actually takes place, it is bound by a known, tolerable, maximum time.

Minsky is credited by Knuth with initiating the development of algorithms for time-sharing garbage collection and list-processing tasks [KNUTH 73]. If two processors are available, these tasks can be performed in parallel, with one of these processors, the collector, responsible for actual garbage collection, and the other performing the list processing and providing the storage requested by a user's program. Dijkstra calls this latter processor the mutator. The collector performs the basic tasks of marking and incorporating unmarked cells to a free list, during which time the mutator is active. The mutator may not request cells until the collector makes them available.

France, Gries, and Muller provide detailed descriptions of Dijkstra's algorithm, but their main concern is to prove correctness [FRANCE 78, GRIES 77, MULLER 76]. An extension of Dijkstra's algorithm with multiple mutators is considered in [IAMPORT 76].

Steele has independently developed a method for parallel garbage collection based on the Minsky-Knuth suggestion. He was one of the first to propose actual algorithms for collecting in parallel. Steele's collector makes exclusive use of semaphores and requires two bits per cell, which are used not only for marking but also for compacting and for readjusting pointers. Compaction is done using the two-pointer technique.

Comparing Dijkstra's to Steele's algorithm is difficult because these authors had different objectives. The former wanted to assure the correctness of his algorithm (regardless of its efficiency), and the latter had in mind an implementation using special hardware, possibly micro-coded.

Kung and Song [KUNG 77] propose a variant of Dijkstra's method that uses four colors for marking and that does not need to trace the free list. The authors prove the correctness of the algorithm and show that it is more efficient than Dijkstra's. To these authors' knowledge, none of the parallel garbage collection algorithms has been implemented, nor are any detailed results from simulation yet available.

An alternative to using two processors is to have one processor time-share the duties of the mutator and collec-

tor. Wadler [WADLER 76] shows that algorithms for performing garbage collection with time-sharing demand a greater percentage of the processing time than does classical sequential garbage collection. This is because the collection effort must proceed even when there is no demand for it.

A second approach for avoiding substantial program interruptions due to garbage collection has been proposed by Baker [see BAKER 78]. His method is an interesting modification of the collector. Baker's modification is such that each time a cell is requested a fixed number of cells, k , are moved from one semispace to the other. This implies that the two semispaces are simultaneously active. In a paging environment, the extra memory required is of less significance than the possible increase in the size of the average working set. Since the moved lists are compacted, page faults are likely to be minimized.

The moving of k cells during a "cons" corresponds to the tracing of many cells in classical garbage collection. By distributing some of the garbage collection tasks during list processing. Baker's method guarantees that actual garbage collection cannot last more than a fixed (tolerable) amount of time: the time to flip the semispaces and to readjust a fixed number of pointers declared in the user's program. Thus his algorithm may be

used in real-time applications

A characteristic of Baker's real-time algorithm is that the size of the semispaces may have to be increased, depending on the value of k and the type of list processing done by the program. In other words, the choice of k expresses the trade-off between the time to execute a cons and the total storage required. For example, for $k = 1/3$, a cell is moved every third time a cons is called. This would speed up the computation but increase the amount of storage required.

In his paper Baker offers an informal proof of his algorithm's correctness and shows how it can be modified to handle varisized cells and arrays of pointers. He also presents analyses of storage requirements of the algorithm and how it compares with those of other garbage collection methods. A IISP machine built at M.I.T used Baker's approach [BAWDEN 77]; its memory is subdivided into areas, and a list of outgoing references is kept for each. Those areas which do not change during program execution are not copied: tracing starts from their corresponding list of outgoing references. This approach, which has been further developed by Bishop [BISHOP 77], is a possible alternative for real-time collection. Another alternative is the use of an auxiliary processor as suggested by Deutsch and Potrow [DEUTSCH 76] in their incremental

garbage collection technique mentioned in the previous section.

3.4. Language Implementation

Recursion is frequently utilized in programs that manipulate linked-list structures. A stack is indispensable for executing these programs. Therefore, separate regions for the allocated cells and for the stack must coexist in the memory. It is true that stacks can be "simulated" by linked lists, so that the memory stores only list structures. However, this is both space and time consuming because an extra field is required to link together the data in the stack and more complex operations are needed for pushing and popping. It is, therefore, easier and more space efficient to implement the stack in contiguous positions of memory.

It has become current practice to divide the available memory into two areas that are allowed to grow from opposite ends. One of these is reserved for a stack using contiguous locations. The other, called the heap, is available to the allocator for providing new cells, also from contiguous locations. With this arrangement, a simple test can be used to trigger garbage collection. When the pointer to the next free stack position meets the

pointer to the next available position in the heap, collection with compaction is invoked to retrieve space for new cells or for stacking. Therefore, the functions "push" and "new" may trigger garbage collection.

In the case of IISP programs, the function "new" corresponds to a "cons", and "push" is used internally by the compiler or interpreter. Collection can be started either by a "cons" or by a stack overflow caused by situations such as great recursion depth or reading long atoms [COHEN 72].

Since the stack is used in implementing recursion, it usually contains pointers to active, useful cells. The marking algorithms of Section 2.1.1. are used to mark not only the structures referred to by the pointer variables of a program but also those structures that are referred to by pointers on the stack. Therefore, means must be provided to recognize whether a stacked quantity is a pointer (tag bits may be used for this purpose) [Deutsch/Scorr-Waite list-traversal technique].

IISP processors sometimes allow a user to invoke the collector. This is useful when he has an idea of the most propitious time for triggering the collection. Also, the function return may be made available to the user. When the free list is used, the returned cells can be immedi-

ately incorporated into the list.

However, when compaction is required (with the heap and stack arrangement), the returned cells may not be available to the allocator until after the next collection. Another problem with the function return is that a cell may be explicitly returned even though there is still a pointer to it. This is sometimes called the dangling reference problem. Thus, care must be taken not to reuse the cell until there are no pointers to it.

Processors for languages like PL/I and PASCAL allow a user to call the function "new" and provide messages when storage is exhausted. The use of the functions "return" or "collect" is implementation dependent. This author is unaware of PASCAL run-time systems that perform fully automatic garbage collection. It is the user's responsibility to keep free lists of unused cells and to check whether a new cell may be obtained from the allocator.

Arnborg [ARNEBERG 73] described the implementation of a SIMULA compiler designed to operate in a virtual memory environment. SIMULA is a language with block structure: variables declared in a block or procedure exist only when the block or procedure is activated. Although it seems at first sight that one could collect the structures referred to by pointer variables upon exiting from the block in

which they are declared, this is not the case. SIMULIA also allows variables of such types as classes, arrays, and tests that may have longer life spans than their originating blocks; if these variables share linked structures with local block variables, collection cannot be done when exiting from a block.

Since Arnborg's proposed implementation operated in a paging environment, one of the objectives of the collection is to reduce the number of page faults. To perform the collection, Arnborg uses a variant of the method proposed by Fenichel and Yochelson [FENICHEL 69] described in Section 3.1.. The Variant can handle varisized cells rather than only simple LISP cells.

A SNOBOL implementation proposed by Hanson [HANSON 77] uses a variation of the garbage collection techniques for collecting varisized cells described in section 2.2.. It is assumed that additional space for the stack can be requested from the operating system, although such requests should be kept to a minimum. An effort is made to reduce collection time by avoiding marking cells which are known to be used throughout the program's execution. For this purpose the heap is subdivided into two areas of consecutive locations: heap 1 and heap 2. The first contains information which is constantly active and never needs to be marked; the second may contain inactive cells

ately incorporated into the list.

However, when compaction is required (with the heap and stack arrangement), the returned cells may not be available to the allocator until after the next collection. Another problem with the function return is that a cell may be explicitly returned even though there is still a pointer to it. This is sometimes called the dangling reference problem. Thus, care must be taken not to reuse the cell until there are no pointers to it.

Processors for languages like PL/I and PASCAL allow a user to call the function "new" and provide messages when storage is exhausted. The use of the functions "return" or "collect" is implementation dependent. This author is unaware of PASCAL run-time systems that perform fully automatic garbage collection. It is the user's responsibility to keep free lists of unused cells and to check whether a new cell may be obtained from the allocator.

Arnborg [ARNEORG 73] described the implementation of a SIMULA compiler designed to operate in a virtual memory environment. SIMULA is a language with block structure: variables declared in a block or procedure exist only when the block or procedure is activated. Although it seems at first sight that one could collect the structures referred to by pointer variables upon exiting from the block in

3.5. Final Remarks

Tables 1-5 summarize the characteristics of the main algorithms described in the corresponding Sections above. The number of references presented in the bibliography bears witness to the importance and interest in garbage collection. In spite of this activity, many facets of garbage collection remain to be investigated. Specifically, very little comparison has been made of the relative efficiencies of many of the algorithms described above.

New developments in hardware are likely to play an important role in speeding up collection. It has already been suggested that new machines should contain extra bits per word to be used for marking, tagging, or counting references. Machines with special hardware for segmentation and list processing have been constructed [PAWDEN 77] and are in experimental operation.

There has been a trend toward designing and implementing collectors for varisized cells stored in large virtual memories. No explicit guidance based on experimental evidence is yet available on how to do this collection efficiently or in real time. Two promising directions, discussed in Section 3.3., involve either using parallel processors or distributing some of the garbage

collection tasks during the actual processing. It is hoped that this will allow the collection to be done within a known, tolerable, maximum time.

Collection in very large virtual memories is another subject that will become increasingly important. The suggested approaches for these collections deserve further study [see FISECP 77].

Table 1: Collecting single-sized cells

| Algorithm | Main references | Auxiliary storage | Mark(M) or tag(T) bit | Comments |
|--|-----------------------|--------------------|-----------------------|---|
| Marking single-sized cells [1] : | | | | |
| Classical | Knuth 73 | Stack | M | Stack may be as large as n [2] |
| Pounded workspace | Knuth 73 (pp. 415) | Limited-size stack | M | Requires more time than the classical algorithm |
| Link reversal (Deutsch/Schorr-Waite) | Schorr 67 Knuth 73 | No stack is needed | M&T | As above (the tag bit may be replaced by the atom bit) |
| Link reversal with bit stack | Wegbreit 72 | Bit stack | M | Similar to the above algorithm |
| Hybrid | Knuth 73 (pp. 502) | Limited-size stack | M&T | uses a combination of stack and link reversal |
| Compaction of single-sized cells: | | | | |
| Two pointer | Knuth 73 | None | None | Only applicable to single-sized cells. Compaction of arbitrary type |
| Moving collectors of single-sized cells: | | | | |
| Minsky | Minsky 67 | None | M | Compaction of linearizing type |
| Fenichel-Yochelson | Fenichel 69 | Stack | None | Uses two semispaces. Compaction of linearizing type |

[1] These marking algorithms may be extended to mark vari-sized cells.
 [2] r is the total number of accessible cells

Table 2: Collecting varisized cells

| Algorithm | Main references | Auxiliary storage | Mark(M) or tag(T) bit | Comments |
|--|-------------------------|--------------------------|-----------------------|--|
| Marking varisized cells: | | | | |
| Classical | Knuth 72 Thorelli 72 | Stack | M | Stack can be stored using an additional field of each cell |
| Link reversal | Thorelli 76 | No stack is | M&T | Variant of Deutsch/schorr-waite link-reversal technique |
| Compactify varisized cells: All these compactors are of the sliding type | | | | |
| Break table | Haddon 67 | None | None | Break table has to be sorted |
| IISP 2 | Knuth 73 (pp. 602) | Additional word per cell | None | Requires three or more scans |
| Moving (Copying) collectors of varisized cells: | | | | |
| Standish | Standish 80 | One field | None | Copies using three passes |

Table 3: Collecting in virtual memory

| Algorithm | Main references | Auxiliary storage | Comments |
|-----------|-----------------|-----------------------------------|---|
| Paker | Paker 76 | None | Uses two semispaces |
| Pishon | Pishon 77 | Space for keeping interarea lists | Designed for use in very large virtual memories |

Table 4: Reference counter (m is the number of available cells)

| Algorithm | Main references | Storage needed | Comments |
|-----------|------------------------|--|--|
| Classical | Collin 60 | An extra field (of size m) per cell | Cannot handle general circular lists |
| Hybrid | Knuth 72 Deutsch 76 | An extra field (of size $h < m$) per cell | Combines reference counters with classical compacting garbage collection |

Table 5: Parallel and Real-Time Collection

| Algorithm | Main reference | Storage needed | Comments |
|------------------------|-------------------|--|--|
| Parallel (Dijkstra) | Dijkstra 76 | No stack & two bits per cell | Main objective is to prove correctness |
| Parallel (Steele) | Steele 75 | Stack, Two bits per cell and several semaphores | Designed to be microcoded. Does compacting |
| Paker | Paker 78 | Two semispaces whose sizes vary at execution time | Moving of accessible cells is done when a new cell is requested |

4. An Implementation Of Garbage Collector In A Network DBMS

When dealing with performance problems, what does the Database Administrator[DEA] do when jobs are aborting because the database is full or out of space? What should the DEA do when complaints start pouring in about response time being slow or jobs running behind schedule?

Space and performance problems in a volatile database are sometimes caused by a high percentage of deleted, but unremoved, records in the database [SEED 81]. The problem of unremoved, deleted records can be solved by garbage collection mechanism.

4.1. Problem Statement

When a delete of a record is enacted under the DBMS, the record is logically deleted from the database, that is, any reference to the deleted item will not be satisfied. However, in some DBMS, the actual physical file record is not deleted. Over time, the DBMS will cease to operate efficiently because of these undeleted records. The problems that develop when deleted records remain unremoved are a shortage of freespace, and a loss of efficiency. Hence, storage will be wasted on inaccessible and useless data. Time will be wasted since the system will be required to do increased I/O. The space allocated will carry the defunct information, and degrade performance.

4.2. Implementation Part Approach

In M.A. Dvorch's system [DVORCH 82], when the deletion of a set owner/member or the deletion of a record is enacted under the DBMS, the set occurrence or the record is logically deleted from the database. However, the actual physical file record is not deleted. Over time, the database will become cluttered with useless information. In answer to this problem, we plan to implement a garbage collection utility to remove the deleted records and to compact this storage.

This enhancement is based on M.A. Dvorch's DBMS in conjunction with W. Stratton's B-Tree program (the logic is original from Wirth and Niklaus E-tree algorithm [see WIRTE 76]), and is written in the C programming language under the UNIX operating system.

In the following sections, we will first of all take a look at the structure of the database files in the implemented DBMS and the functions of Database deletion that give the background of the garbage collector.

4.2.1. Database Files

Figure 4-1 depicts the implemented system on which this project is based [DVONCH 82]. The Database Management system presently implemented provides (1) definition directives, which allow the creation of record and set types, and (2) manipulation commands or primitives, which allow the manipulation of the data stored in the database.

This model uses two different form types to represent the information contained in the database. A record type is the format of the record. A set type is the definition of a relationship between two record types. For a set type, one record type is the owner type and another record type is the member type.

Definition directives are accepted by the DBMS and appropriate record or set types are defined. The DBMS does not allow the user to delete a record type or set type once it has been defined. The record type and set type definition are maintained in a binary tree in memory while the DBMS is in use. When a new record or set type is defined, a node is created for it in the appropriate tree. At the conclusion of the database update, pertinent information from the trees is stored in a file named dbm_file which is the data base management file as shown in Figure 4-1. The first field in the dbm_file is the

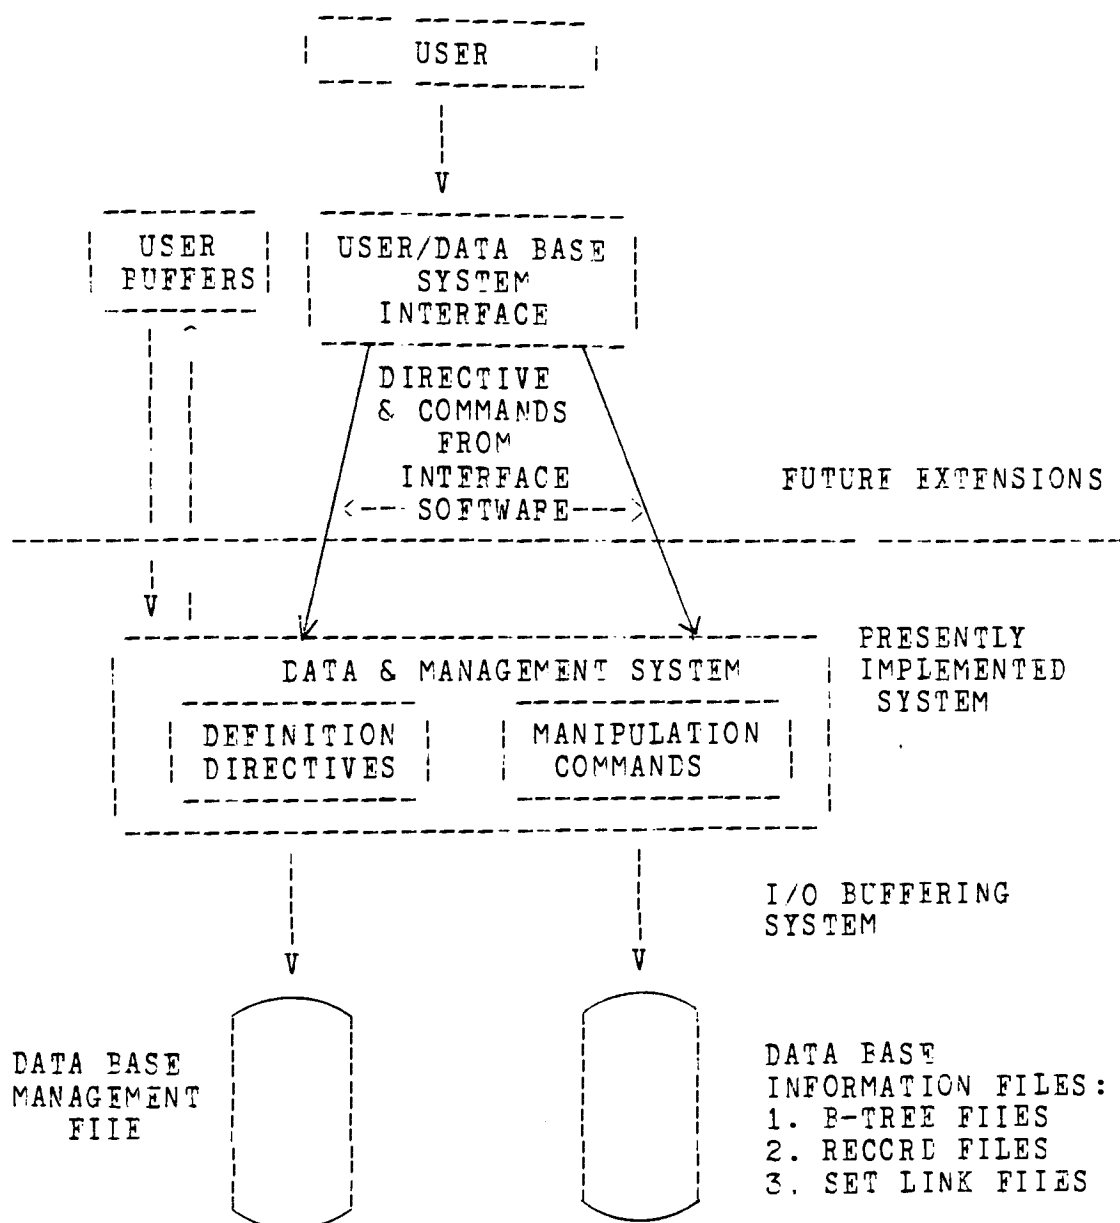


FIGURE 4-1
SYSTEM ON WHICH THE PROJECT IS BASED

number of record types defined within the database, the second is the number of set types defined. After these two fields come the record type information and then the set type information. When the database management system is run, the system builds the new binary type trees from the information stored on the dbm_file. The dbm_file is a normal UNIX file that can be examined by and used with other UNIX tools. If no record types or set types have been defined, the system will print a message informing the user of this fact.

4.2.1.1. Record Types

The system creates the record type by inserting the record type name with ".rf" appended to it into the binary tree containing record type definitions. When the user then issues an "add record" command for a particular record type, the database system fills an input/output buffer with the desired record, checks the field count on the record for an error, and builds the record's key. The key is also error checked for duplicity. The record is then passed to the input/output buffering system where it is stored at the end of a file whose name is the record type name with the suffix ".rf". As the record is passed to the buffering system, the location of the record within

the file is determined. The location of a record is the block offset from the beginning of the file and then an offset to the position within the block. Once the address of the record within the file is known, the address information and the record key is stored in the record type B-tree named "dbm.rec.bt". If the record type participates as an owner in any set type, the new record is added as an owner occurrence for each such set. These records are stored in regular UNIX file. Again, this allows the stored database files to be used with any other UNIX tools. For example, a record file could be printed directly. This separation of record files from database information is one of the strengths of the system, for it is in keeping with the UNIX idea of software tools [M.A. Dvornch's thesis 82]. However, deleted records still remain on the record file until garbage collection is done.

4.2.1.2. Set Types

The system creates a set type by inserting the set type name with ".sf" appended to it into the set type binary tree containing set type definition.

After the set type has been defined, owner occurrences are automatically created whenever a record of

the owner record type is added to the database. When an owner record occurrence is added to the database, a link-record for the owner is built. This link-record contains pointers with the forward pointer first and backward pointer second. The front pointer of any owner or member in a set occurrence contains the key of the next owner or member in the set occurrence. The back pointer of any member contains the key of the prior member or the owner. The owner's back pointer is always null. This link-record is then given to the I/O buffer system to be added to a file named by the set type name with "sf" appended to it. As the record is passed to the I/O buffer system, information is gathered on the location of the pointer record within the set type file. This location information and the owner key are stored in the set type B-tree. The link-records are stored as regular UNIX files and, therefore, can be used as arguments for other UNIX tools.


If the user wishes to add a member to a set occurrence, he must give an explicit "add member" command. The user specifies the set type and owner key of the appropriate set occurrence. Inserting the new member into a set occurrence is accomplished by building a link-record for the new member. These link-records make up a chain that connect the set occurrences, and can be traversed to obtain all the members of a set occurrence. The new

member occurrence is added as the immediate successor to the owner. Order of retrieval is therefore LIFO (Figure 1-2). Once the link-record is stored, the file offset information along with the new member key are stored in the set type B-tree.

4.2.2. Database Deletion

The database management system's deletion tools are recursive. "Since deleting a record implies possible deletion of set occurrences and vice versa, recursion is required. Each deletion tool was designed to do its own

ORIGINAL SET:



AFTER ADDITION OF MEMBER 3:




FIGURE 4-2
ADDITION OF NEW MEMBER TO SET OCCURRENCE

particular job and then call a friend to do the rest"
[Dvonch 92].

4.2.2.1. Delete Record

When a record is deleted from the database, the record key and file offset are removed from the record type B-tree. Since the offsets for the actual record can no longer be referenced, the record is logically deleted from the database. However, the actual physical record will still remain on the record type file. This record then becomes the province of the garbage collection system.

Once the record is logically deleted, delete record recursively searches the set type binary tree looking for sets in which this record type participates as an owner or member. If a set is found, delete owner or delete member is called to remove the discarded record key from the set.

4.2.2.2. Delete Member

The delete member tool causes a member's key and link-record offset to be removed from the set type B-tree. The pointers in the link-records of prior and post set occurrence members are reset to exclude the deleted

member. The actual link-record is left on the set type file, requiring garbage collection at a later time. However, since this link-record can no longer be referenced, it is effectively deleted from the database.

4.2.2.3. Delete Owner

When the delete owner command is given, the owner's key and link-record file offset are removed from the set type F-tree. Next, "delete member" is repeatedly called until all of this owner's members are deleted. Finally, "delete record" is called to delete the owner's record from the record type file. That is, when an owner is deleted from a set, each of his members is also deleted.

4.2.3. Prototype

Consider a simple prototype database which consists of four record types and three set types with certain relationships expressed through set occurrences (see Figure 4-3). And using the same test data files run by M.A. Dvorch's DBMS as the test data files of this garbage collector's implementation project for consistency.

Four record type descriptions and contents:

(1)

Record Type Name: faculty
Field Delimiter: *
Number of Fields: 5
Number of Key Fields: 1
Position of Key Fields: 2

Contents:

Peter*A1*10*A186*25
Bill*A2*10*2132*57
Roy*3A*10*A285*72
Jack*4A*10*1116*13

(2)

Record Type Name: student
Field Delimiter: :
Number of Fields: 4
Number of Key Fields: 1
Position of Key Fields: 3

Contents:

Mary:CAST:E1:Comp Scie
Leslie:CAST:E2:Comp Scie
John:SP:3B:PPPD
Tom:CAST:4B:Syst Soft
Mary:SP:5B:PPPD

(3)

Record Type Name: housing
Field Delimiter: *
Number of Fields: 3
Number of Key Fields: 1
Position of Key Fields: 1

Contents:

405*Billings*25
215*Watson*1105

(4)

Record Type Name: courses
Field Delimiter: *
Number of Fields: 8
Number of Key Fields: 4
Position of Key Fields: 5 1 3 4

Contents:

P1*0601*81*1*875*1*D*r
P2*0601*81*2*875*1*A*nr
P1*0601*81*1*720*2*F*nr
3P*0532*81*2*875*1*A*nr
B2*0532*81*2*859*1*F*r
B1*0601*81*2*875*1*F*r
5P*0601*80*2*875*1*D*r
5P*0601*81*3*875*1*F*r

Three set type descriptions and contents:

(1)

Set Type Name: fs
 Owner Record Type: faculty
 Member Record Type: student
 Relationship: the faculty member is the advisor
 of the student

Contents:

A1 owns
 P1
 3P

A2 owns
 P2
 4P

3A owns
 5P

4A has no members

(2)

Set Type Name: sc
 Owner Record Type: student
 Member Record Type: courses
 Relationship: the student has taken the course

Contents:

P1 owns
 875P1811
 720P1811
 875P1812

P2 owns
 875P2812
 850b2812

3P owns
 8753B812

4F has no members

5F owns
8755B802
8755B813

(3)

Set Type Name: hs
Owner Record Type: housing
Member Record Type: student
Relationship: the student is a resident of
the housing complex

Contents:

405 owns
F1
5E


216 owns
3E

To delete student P1's record in this database, we use the delete record command ("dr student P1"). (see Figure 4-4).

The "dr student P1" command would remove the record key P1 and file offset from the record type P-tree. Then P1 would be removed from the faculty-student set (file: fs.sf) and the housing-student set (file: hs.sf) as a member (call "dm fs P1" and "dm hs P1" respectively) and the student-courses set (file: sc.sf) as an owner (call "do sc P1") .

The "do sc P1" command would remove P1 from the sc set. Next, courses 875P1811, 720P1811, and 875P1812 would be removed from the sc set (call "dr sc 875P1811", "dm sc 720P1811", "dm sc 875P1812") and the courses record file (call "dr courses 875P1811", "dr courses 720P1811", "dr courses 875P1812") and then P1 would be removed from the student record file.

Figure 4-4 depicts the Database files when the DPMS is not being used after executing the delete command "dr student P1". As we can see, it needs garbage collection to recognize the database information files, ie record files and set link files which contains the useless records after deletion.



| DATABASE MANAGEMENT FILE | | DATABASE INFORMATION FILES | | 2. | SET LINK FILES | 3. | B-TREE FILES |
|--------------------------|--|----------------------------|----------------|--|-----------------------|-------------------------|--|
| dbm_file | | RECORD FILES | SET LINK FILES | RECORD TYPE B-TREE FILE | SET TYPE B-TREE FILES | RECORD TYPE B-TREE FILE | SET TYPE B-TREE FILES |
| 4 2 | | faculty.rf | fs.sf | AI BI 3B A2 B2 4B 3A CB 4A | fs.sf | RAT "dbm.rec.bt" | LOC. A1 LOC. A2 LOC. 3A LOC. 4A LOC. B1 LOC. B2 LOC. 3E LOC. 4B LOC. 5B LOC. 40S LOC. 216 LOC. 875E2B12 LOC. 875E2B12 LOC. 875E2B12 LOC. 875E2B12 LOC. 875E2B12 LOC. 875E2B12 LOC. 875E2B12 |
| 5 | | Student.rf | SC.SF | BI 3B 4B 5B 6B 7B 8B 9B 10B 11B 12B 13B 14B 15B 16B 17B 18B 19B 20B 21B 22B 23B 24B 25B 26B 27B 28B 29B 30B 31B 32B 33B 34B 35B 36B 37B 38B 39B 40B 41B 42B 43B 44B 45B 46B 47B 48B 49B 50B 51B 52B 53B 54B 55B 56B 57B 58B 59B 60B 61B 62B 63B 64B 65B 66B 67B 68B 69B 70B 71B 72B 73B 74B 75B 76B 77B 78B 79B 80B 81B 82B 83B 84B 85B 86B 87B 88B 89B 90B 91B 92B 93B 94B 95B 96B 97B 98B 99B 100B 101B 102B 103B 104B 105B 106B 107B 108B 109B 110B 111B 112B 113B 114B 115B 116B 117B 118B 119B 120B 121B 122B 123B 124B 125B 126B 127B 128B 129B 130B 131B 132B 133B 134B 135B 136B 137B 138B 139B 140B 141B 142B 143B 144B 145B 146B 147B 148B 149B 150B 151B 152B 153B 154B 155B 156B 157B 158B 159B 160B 161B 162B 163B 164B 165B 166B 167B 168B 169B 170B 171B 172B 173B 174B 175B 176B 177B 178B 179B 180B 181B 182B 183B 184B 185B 186B 187B 188B 189B 190B 191B 192B 193B 194B 195B 196B 197B 198B 199B 200B 201B 202B 203B 204B 205B 206B 207B 208B 209B 210B 211B 212B 213B 214B 215B 216B 217B 218B 219B 220B 221B 222B 223B 224B 225B 226B 227B 228B 229B 230B 231B 232B 233B 234B 235B 236B 237B 238B 239B 240B 241B 242B 243B 244B 245B 246B 247B 248B 249B 250B 251B 252B 253B 254B 255B 256B 257B 258B 259B 260B 261B 262B 263B 264B 265B 266B 267B 268B 269B 270B 271B 272B 273B 274B 275B 276B 277B 278B 279B 280B 281B 282B 283B 284B 285B 286B 287B 288B 289B 290B 291B 292B 293B 294B 295B 296B 297B 298B 299B 300B 301B 302B 303B 304B 305B 306B 307B 308B 309B 310B 311B 312B 313B 314B 315B 316B 317B 318B 319B 320B 321B 322B 323B 324B 325B 326B 327B 328B 329B 330B 331B 332B 333B 334B 335B 336B 337B 338B 339B 340B 341B 342B 343B 344B 345B 346B 347B 348B 349B 350B 351B 352B 353B 354B 355B 356B 357B 358B 359B 360B 361B 362B 363B 364B 365B 366B 367B 368B 369B 370B 371B 372B 373B 374B 375B 376B 377B 378B 379B 380B 381B 382B 383B 384B 385B 386B 387B 388B 389B 390B 391B 392B 393B 394B 395B 396B 397B 398B 399B 400B 401B 402B 403B 404B 405B 406B 407B 408B 409B 410B 411B 412B 413B 414B 415B 416B 417B 418B 419B 420B 421B 422B 423B 424B 425B 426B 427B 428B 429B 430B 431B 432B 433B 434B 435B 436B 437B 438B 439B 440B 441B 442B 443B 444B 445B 446B 447B 448B 449B 450B 451B 452B 453B 454B 455B 456B 457B 458B 459B 460B 461B 462B 463B 464B 465B 466B 467B 468B 469B 470B 471B 472B 473B 474B 475B 476B 477B 478B 479B 480B 481B 482B 483B 484B 485B 486B 487B 488B 489B 490B 491B 492B 493B 494B 495B 496B 497B 498B 499B 500B 501B 502B 503B 504B 505B 506B 507B 508B 509B 510B 511B 512B 513B 514B 515B 516B 517B 518B 519B 520B 521B 522B 523B 524B 525B 526B 527B 528B 529B 530B 531B 532B 533B 534B 535B 536B 537B 538B 539B 540B 541B 542B 543B 544B 545B 546B 547B 548B 549B 550B 551B 552B 553B 554B 555B 556B 557B 558B 559B 560B 561B 562B 563B 564B 565B 566B 567B 568B 569B 570B 571B 572B 573B 574B 575B 576B 577B 578B 579B 580B 581B 582B 583B 584B 585B 586B 587B 588B 589B 590B 591B 592B 593B 594B 595B 596B 597B 598B 599B 600B 601B 602B 603B 604B 605B 606B 607B 608B 609B 610B 611B 612B 613B 614B 615B 616B 617B 618B 619B 620B 621B 622B 623B 624B 625B 626B 627B 628B 629B 630B 631B 632B 633B 634B 635B 636B 637B 638B 639B 640B 641B 642B 643B 644B 645B 646B 647B 648B 649B 650B 651B 652B 653B 654B 655B 656B 657B 658B 659B 660B 661B 662B 663B 664B 665B 666B 667B 668B 669B 670B 671B 672B 673B 674B 675B 676B 677B 678B 679B 680B 681B 682B 683B 684B 685B 686B 687B 688B 689B 690B 691B 692B 693B 694B 695B 696B 697B 698B 699B 700B 701B 702B 703B 704B 705B 706B 707B 708B 709B 710B 711B 712B 713B 714B 715B 716B 717B 718B 719B 720B 721B 722B 723B 724B 725B 726B 727B 728B 729B 730B 731B 732B 733B 734B 735B 736B 737B 738B 739B 740B 741B 742B 743B 744B 745B 746B 747B 748B 749B 750B 751B 752B 753B 754B 755B 756B 757B 758B 759B 760B 761B 762B 763B 764B 765B 766B 767B 768B 769B 770B 771B 772B 773B 774B 775B 776B 777B 778B 779B 780B 781B 782B 783B 784B 785B 786B 787B 788B 789B 790B 791B 792B 793B 794B 795B 796B 797B 798B 799B 800B 801B 802B 803B 804B 805B 806B 807B 808B 809B 810B 811B 812B 813B 814B 815B 816B 817B 818B 819B 820B 821B 822B 823B 824B 825B 826B 827B 828B 829B 830B 831B 832B 833B 834B 835B 836B 837B 838B 839B 840B 841B 842B 843B 844B 845B 846B 847B 848B 849B 850B 851B 852B 853B 854B 855B 856B 857B 858B 859B 860B 861B 862B 863B 864B 865B 866B 867B 868B 869B 870B 871B 872B 873B 874B 875B 876B 877B 878B 879B 880B 881B 882B 883B 884B 885B 886B 887B 888B 889B 890B 891B 892B 893B 894B 895B 896B 897B 898B 899B 900B 901B 902B 903B 904B 905B 906B 907B 908B 909B 910B 911B 912B 913B 914B 915B 916B 917B 918B 919B 920B 921B 922B 923B 924B 925B 926B 927B 928B 929B 930B 931B 932B 933B 934B 935B 936B 937B 938B 939B 940B 941B 942B 943B 944B 945B 946B 947B 948B 949B 950B 951B 952B 953B 954B 955B 956B 957B 958B 959B 960B 961B 962B 963B 964B 965B 966B 967B 968B 969B 970B 971B 972B 973B 974B 975B 976B 977B 978B 979B 980B 981B 982B 983B 984B 985B 986B 987B 988B 989B 990B 991B 992B 993B 994B 995B 996B 997B 998B 999B 1000B 1001B 1002B 1003B 1004B 1005B 1006B 1007B 1008B 1009B 1010B 1011B 1012B 1013B 1014B 1015B 1016B 1017B 1018B 1019B 1020B 1021B 1022B 1023B 1024B 1025B 1026B 1027B 1028B 1029B 1030B 1031B 1032B 1033B 1034B 1035B 1036B 1037B 1038B 1039B 1040B 1041B 1042B 1043B 1044B 1045B 1046B 1047B 1048B 1049B 1050B 1051B 1052B 1053B 1054B 1055B 1056B 1057B 1058B 1059B 1060B 1061B 1062B 1063B 1064B 1065B 1066B 1067B 1068B 1069B 1070B 1071B 1072B 1073B 1074B 1075B 1076B 1077B 1078B 1079B 1080B 1081B 1082B 1083B 1084B 1085B 1086B 1087B 1088B 1089B 1090B 1091B 1092B 1093B 1094B 1095B 1096B 1097B 1098B 1099B 1100B 1101B 1102B 1103B 1104B 1105B 1106B 1107B 1108B 1109B 1110B 1111B 1112B 1113B 1114B 1115B 1116B 1117B 1118B 1119B 1120B 1121B 1122B 1123B 1124B 1125B 1126B 1127B 1128B 1129B 1130B 1131B 1132B 1133B 1134B 1135B 1136B 1137B 1138B 1139B 1140B 1141B 1142B 1143B 1144B 1145B 1146B 1147B 1148B 1149B 1150B 1151B 1152B 1153B 1154B 1155B 1156B 1157B 1158B 1159B 1160B 1161B 1162B 1163B 1164B 1165B 1166B 1167B 1168B 1169B 1170B 1171B 1172B 1173B 1174B 1175B 1176B 1177B 1178B 1179B 1180B 1181B 1182B 1183B 1184B 1185B 1186B 1187B 1188B 1189B 1190B 1191B 1192B 1193B 1194B 1195B 1196B 1197B 1198B 1199B 1200B 1201B 1202B 1203B 1204B 1205B 1206B 1207B 1208B 1209B 1210B 1211B 1212B 1213B 1214B 1215B 1216B 1217B 1218B 1219B 1220B 1221B 1222B 1223B 1224B 1225B 1226B 1227B 1228B 1229B 1230B 1231B 1232B 1233B 1234B 1235B 1236B 1237B 1238B 1239B 1240B 1241B 1242B 1243B 1244B 1245B 1246B 1247B 1248B 1249B 1250B 1251B 1252B 1253B 1254B 1255B 1256B 1257B 1258B 1259B 1260B 1261B 1262B 1263B 1264B 1265B 1266B 1267B 1268B 1269B 1270B 1271B 1272B 1273B 1274B 1275B 1276B 1277B 1278B 1279B 1280B 1281B 1282B 1283B 1284B 1285B 1286B 1287B 1288B 1289B 1290B 1291B 1292B 1293B 1294B 1295B 1296B 1297B 1298B 1299B 1300B 1301B 1302B 1303B 1304B 1305B 1306B 1307B 1308B 1309B 1310B 1311B 1312B 1313B 1314B 1315B 1316B 1317B 1318B 1319B 1320B 1321B 1322B 1323B 1324B 1325B 1326B 1327B 1328B 1329B 1330B 1331B 1332B 1333B 1334B 1335B 1336B 1337B 1338B 1339B 1340B 1341B 1342B 1343B 1344B 1345B 1346B 1347B 1348B 1349B 1350B 1351B 1352B 1353B 1354B 1355B 1356B 1357B 1358B 1359B 1360B 1361B 1362B 1363B 1364B 1365B 1366B 1367B 1368B 1369B 1370B 1371B 1372B 1373B 1374B 1375B 1376B 1377B 1378B 1379B 1380B 1381B 1382B 1383B 1384B 1385B 1386B 1387B 1388B 1389B 1390B 1391B 1392B 1393B 1394B 1395B 1396B 1397B 1398B 1399B 1400B 1401B 1402B 1403B 1404B 1405B 1406B 1407B 1408B 1409B 1410B 1411B 1412B 1413B 1414B 1415B 1416B 1417B 1418B 1419B 1420B 1421B 1422B 1423B 1424B 1425B 1426B 1427B 1428B 1429B 1430B 1431B 1432B 1433B 1434B 1435B 1436B 1437B 1438B 1439B 1440B 1441B 1442B 1443B 1444B 1445B 1446B 1447B 1448B 1449B 1450B 1451B 1452B 1453B 1454B 1455B 1456B 1457B 1458B 1459B 1460B 1461B 1462B 1463B 1464B 1465B 1466B 1467B 1468B 1469B 1470B 1471B 1472B 1473B 1474B 1475B 1476B 1477B 1478B 1479B 1480B 1481B 1482B 1483B 1484B 1485B 1486B 1487B 1488B 1489B 1490B 1491B 1492B 1493B 1494B 1495B 1496B 1497B 1498B 1499B 1500B 1501B 1502B 1503B 1504B 1505B 1506B 1507B 1508B 1509B 1510B 1511B 1512B 1513B 1514B 1515B 1516B 1517B 1518B 1519B 1520B 1521B 1522B 1523B 1524B 1525B 1526B 1527B 1528B 1529B 1530B 1531B 1532B 1533B 1534B 1535B 1536B 1537B 1538B 1539B 1540B 1541B 1542B 1543B 1544B 1545B 1546B 1547B 1548B 1549B 1550B 1551B 1552B 1553B 1554B 1555B 1556B 1557B 1558B 1559B 1560B 1561B 1562B 1563B 1564B 1565B 1566B 1567B 1568B 1569B 1570B 1571B 1572B 1573B 1574B 1575B 1576B 1577B 1578B 1579B 1580B 1581B 1582B 1583B 1584B 1585B 1586B 1587B 1588B 1589B 1590B 1591B 1592B 1593B 1594B 1595B 1596B 1597B 1598B 1599B 1600B 1601B 1602B 1603B 1604B 1605B 1606B 1607B 1608B 1609B 1610B 1611B 1612B 1613B 1614B 1615B 1616B 1617B 1618B 1619B 1620B 1621B 1622B 1623B 1624B 1625B 1626B 1627B 1628B 1629B 1630B 1631B 1632B 1633B 1634B 1635B 1636B 1637B 1638B 1639B 1640B 1641B 1642B 1643B 1644B 1645B 1646B 1647B 1648B 1649B 1650B 1651B 1652B 1653B 1654B 1655B 1656B 1657B 1658B 1659B 1660B 1661B 1662B 1663B 1664B 1665B 1666B 1667B 1668B 1669B 1670B 1671B 1672B 1673B 1674B 1675B 1676B 1677B 1678B 1679B 1680B 1681B 1682B 1683B 1684B 1685B 1686B 1687B 1688B 1689B 1690B 1691B 1692B 1693B 1694B 1695B 1696B 1697B 1698B 1699B 1700B 1701B 1702B 1703B 1704B 1705B 1706B 1707B 1708B 1709B 1710B 1711B 1712B 1713B 1714B 1715B 1716B 1717B 1718B 1719B 1720B 1721B 1722B 1723B 1724B 1725B 1726B 1727B 1728B 1729B 1730B 1731B 1732B 1733B 1734B 1735B 1736B 1737B 1738B 1739B 1740B 1741B 1742B 1743B 1744B 1745B 1746B 1747B 1748B 1749B 1750B 1751B 1752B 1753B 1754B 1755B 1756B 1757B 1758B 1759B 1760B 1761B 1762B 1763B 1764B 1765B 1766B 1767B 1768B 1769B 1770B 1771B 1772B 1773B 1774B 1775B 1776B 1777B 1778B 1779B 1780B 1781B 1782B 1783B 1784B 1785B 1786B 1787B 1788B 1789B 1790B 1791B 1792B 1793B 1794B 1795B 1796B 1797B 1798B 1799B 1800B 1801B 1802B 1803B 1804B 1805B 1806B 1807B 1808B 1809B 1810B 1811B 1812B 1813B 1814B 1815B 1816B 1817B 1818B 1819B 1820B 1821B 1822B 1823B 1824B 1825B 1826B 1827B 1828B 1829B 1830B 1831B 1832B 1833B | | | |

| | |
|--------------|-------------------------------------|
| MANIPULATION | q ar dr fr co |
| COMMANDS | gr ao do fo ca gs am dm ff fn |

INFORMATION FILES

| | | | | | | |
|----|--------------------|---------|---|---|---|-----|
| 13 | faculty | * 5 | 1 | 2 | 0 | 76 |
| | courses | * 8 | 4 | 5 | 1 | 34 |
| | student | : 4 | 1 | 2 | 0 | 122 |
| | housing | * 3 | 1 | 1 | 0 | 32 |
| | is faculty student | 0 | | | | 387 |
| | sc student | courses | 0 | | | 559 |
| | is housing student | 0 | | | | 215 |

faculty.rf
Peter*A1*10*A186*25
Bill*A2*10*2132*57
Roy*A3*10*A285*72
Jack*A4*10*1116*13

Student, rf

Mary:CAST:B1:Comp Scie
 Leslie:CAST:B2:Comp Scie
 John:SP:3B:PPPD
 Tom:CAST:4B:syst Soft
 Mary:SP:5B:PPPD

housing.rf

405*Billings*25
216*Watson*1105

courses.rf

[illegible]

45.5

[illegible]

fs.sc

[illegible]

fs.sf

[illegible]

2. SET LINK FILES

3. B-TREE FILES

RECORD TYPE B-TREE FILE

RBT " dbm.rec.bt "

10C. A1
10C. A2
10C. 3A
10C. 4A
10C. B1
10C. B2
10C. 3B
10C. 4B
10C. 5B
10C. 4C
10C. 216
10C. 875
10C. 875
10C. 720
10C. 875
10C. 850
10C. 875
10C. 875
10C. 875

SET TYPE B-TREE FILES

79.35

| | |
|------------------|------------|
| loc. B1(o) | loc. A1(o) |
| loc. 875B1811(m) | loc. B1(m) |
| loc. 725B1811(m) | loc. 3B(m) |
| loc. 675B1812(m) | loc. A2(o) |
| loc. B2(o) | loc. B2(m) |
| loc. 675B2812(m) | loc. 4B(m) |
| loc. 650B2812(m) | loc. 3A(o) |
| loc. 3B(c) | loc. 5B(m) |
| loc. 675B2812(m) | loc. 4A(o) |

hs. 6t

loc. 405(o)
loc. B1(m)
loc. 5b(n)
loc. 216(o)
loc. 3b(n)

FIGURE 4-3

DATABASE FILES
 AFTER THE DBMS
 HAS BEEN RUN.

4.2.4. Algorithms

For the system involved, the approach to the garbage collection problem for record type files is the inspect and slide function. Armed with the record type definition, the garbage collector moves down the record file, inspecting each record. The key for the record is recreated using the record type definition. Then the P-tree is searched for that key. If the key is present in the P-tree, the record slides as far forward on the file as possible, and then the offsets of this record in the P-tree reset to reflect the new location. If the key was not in the P-tree, the record overlaid when a valid record from farther down the file was slid over it. This method works well for the record type files; however, it can not work for the set type files that contain link pointers information.

The method used for the set type files is a P-tree walk and move procedure. A tool is built that would visit each node in the P-tree, the set link record is moved to a new file and the offsets readjusted to fit the link record's new placement. Once the new "clean" file is built, the old file space is returned to the system.

Here, two utilities: "gr" (record garbage collector) and "gs" (set garbage collector) are implemented. There

are many ways to collect wasted file space. The designer of the garbage collection tool will have to weigh the merits of each possibility, and implement the most reasonable approach.

BIBLIOGRAPHY

Arnborg, S.. Optimal memory management in a system with garbage collection, FIT 14, 4(1974), pp. 375-381

Paker, H.G., List-processing in real time on a serial computer CACM 21, 4(april 1978), pp. 280-294

Parth, J.M.. Shifting garbage collection overhead to compile time CACM 20, 7(July 1977), pp. 513-518

Pawder, A., Greenblatt, R., Holloway, J., Knight, T., Moon, D., and Weinreb, D., Lisp machine progress report, Memo 4444, A.I. Lab. M.I.T., Cambridge, Mass., Aug. 1977

Fishop, P.P., Computer systems with a very large address space and garbage collection Lab. for Computer Science, MIT Rep., TR-178, M.I.T., Cambridge, Mass., May 1977

Pobrow, D.G., and Murphy, D.L. structure of a LISP system using two-level storage CACM 10, 3(March 1967), pp. 155-159

Pobrow, D.G., and Murphy, D.L., A note on the efficiency of a LISP computation in a paged machine, CACM 11, 8(Aug. 1968), pp. 556-560

Pobrow, D.G., A note on hash linking CACM 18, 7(July 1975), pp. 413-415

Cherey, C.J., A nonrecursive list compacting algorithm
CACM 13, 11(Nov. 1970), pp. 677-678

Clark, D.W., A fast algorithm for copying binary trees,
Inf. Process. Lett. 9, 3(Dec. 1975), pp. 62-63

Clark, D.W., An efficient list moving algorithm using constant workspace, CACM 19, 6(June 1976), pp. 352-354

Clark, D.W., and Green, C.C., An empirical study of list structure in Lisp, CACM 20, 2(Feb. 1977), pp. 78-86

Clark, D.W., A fast algorithm for copying list structures,
CACM 21, 5(May 1978), pp. 351-357

Cohen, J., and Trilling, L., Remarks on Garbage Collection using a two level storage, BIT 7, 1(Jan. 1967), pp. 22-30

Cohen, J., Use of fast and slow memories in list-processing languages CACM 10, 2(Feb. 1967), pp. 82-86

Cohen, J., and Zuckerman, C., Evaluate in simple Fortran: A tutorial on interpreting Lisp, BIT 12, 3(1972), pp. 299-317

Collin, B.F., A method for overlapping and erasure of lists, CACAM 3, 12(Dec. 1960), pp. 655-657

Date, C. J., An Introduction To Database Systems,
Addison-Wasley Publishing Company, 1977

Deitel, H. M., An Introduction to Operation System
Addison-Wasley Publishing Company, 1983

Deutsch, L.P., An efficient incremental automatic garbage collector CACM 19. 9(Sept. 1976), pp. 522-526

Dijkstra, E.W., Lamport, L., Martin, A.J., Scholten, C.S.,
and Steffens, E.F.M., On-the-fly garbage collection: An exercise in cooperation, CACM 21. 11(Nov 1978), pp. 966-975

Dvorch, M.A., The Implementation Of A Network Oriented Database Management System Designed To Run Under The UNIX Operation System RIT thesis, Jan,1982.

Fenichel, E.R. and Yochelson, J.C., A IISF Garbage-Collector for Virtual-Memory Computer Systems, CACM 12,11(Nov,1969) pp. 611-612

Fisher, D.A., Bounded workspace garbage collection in an address order preserving list processing environment, Inf. Process. Lett., 3. 1(July 1974). pp. 29-32

Fisher, D.A., Copying cyclic list structure in linear time using bounded workspace, CACM 18,5(May 1975), pp. 251-252

Fitch, J.P., Norman, A.C. a note on compacting garbage collection, Comput. J. 21. 1(Feb. 1978), pp. 31-34

France, N., An application of a method for analysis of

cyclic programs, IEEE Trans. Softw. Eng. 4, 5(Sept. 1978), pp. 371-377

Fredman, D.F., and Wise, D.S., Garbage collecting a heap which included a scatter table, Inf. Process Lett. 5, 6(Dec. 1976), pp. 161-164

Gotlieb, C.C., and Gotlieb, L.R., Data types and structures, Prentice-Hall, Englewood Cliffs, N.J., 1978

Gries, D., An exercise in proving parallel programs correct, CACM 20, 12(Dec. 1977), pp. 921-930

Haddon, B.K., and Waite, W.M., A compaction procedure for variable length storage element. Comput. J., 10(Aug. 1967), pp. 162-165

Hansen, W.J., Compact List Representation: Definition, Garbage Collection, And System Implementation CACM 12,9(Sep,1969), pp. 499-507

Hanson, D.R., Storage management for an implementation of Snobol 4, Software: Practice and Experience 7, 2(1977), pp. 179-192

Horowitz, E., Sahni, S., Fundamentals Of Data Structures, Computer Science Press, Inc. 1976

Kernighan and Ritchie, The C Programming Language, Prentice-Hall, INC., 1978

Kernighan and Ritchie, "UNIX Programming - Second Edition"

Knuth, D. E., The Art of Computer Programming, Vol.1, Fundamental Algorithms Addison-Wasley Publishing Company, 1973

Lamport, L., Garbage collection with multiple processes: An exercise in parallelism, Proc. IEEE conf. Parallel Processing, Aug. 1976

Lang, P., and Wegbreit, B., Fast compactification, Rep. 25-72, Harvard Univ., Cambridge, Mass., Nov. 1972

Martin J., Computer Data-Base Organization Prentice-Hall, INC. 1977

Marshall, S., An Algol-68 garbage collector, in Algol 68 implementation J.E.I. peck (Ed.), North-Holland, Amsterdam, 1971, pp. 239-243

Minsky, M.L., A Lisp garbage collector algorithm using serial secondary storage Memo 58 (rev.), Project MAC, M.I.T., Cambridge, Mass, Dec. 1963

Moon, D.A., MACLisp reference manual, Project MAC, M.I.T., Cambridge, Mass, April 1974

Morris, F.L., A time- and space- efficient garbage collection algorithm, CACM 21, 8(Aug. 1978), pp. 662-665

Morris, F.L., On a comparison of garbage collection tech-

niques, technical correspondence, CACM 22, 10(Oct. 1979), pp. 571

Muller, K.G., On the feasibility of concurrent garbage collection Ph.D thesis, Tech. Hogeschool Delft, March 1976

Pfaltz, J.L., Computer Data Structures McGraw-Hill, Inc. 1977

Reingold, E.M., A nonrecursive list moving algorithm, CACM 16, 5(May 1973), pp. 305-307

Robson, J.M., A bounded storage algorithm for copying cyclic structures, CACM 20, 6(June 1977), pp. 431-433

Ross, D.T., The AED free storage package CACM 10, 8(Aug. 1967), pp. 481-492

Schorr, H., Waite, M., ul 2 A Efficient Machine-Independent Procedure For Garbage Collection Various List Structures, CACM 10,8(Aug,1967), pp. 501-506

Schorr, H., and Yochelson, J.C., An Efficient Machine-Independent Procedure For Garbage Collection In Various List Structures, CACM 10,8(Aug,1967), pp. 501-506

SEED---a course for new users International DP system Inc., 1981.

Standish, T.A . Data structures techniques Addison-wesley,

Reading, Mass., 1980

Steele, G.L., Multiprocessing Compactifying Garbage Collection, CACM 19, 9(Sept. 1976), pp. 498-528

Terashima, M., and Goto, E., Genetic order and compactifying garbage collections, Inf. Process. Lett. 7, 1(Jan. 1978), pp. 27-32

Thorelli, L.E., A fast compactifying garbage collector, BIT 16, 4(1976), pp. 426-441

Wadler, P.L., Analysis Of An Algorithm For Real Time Garbage Collection, CACM 19,9(Sept. 1976), pp. 491-500

Waite, W.M., Implementing software for non-numeric applications, Prentice-Hall, englewood cliffs, N.J., 1973

Walden, D.C., A note on Cheney's nonrecursive list-compacting algorithm, CACM 15, 4(April 1972), pp. 275

Wegbreit, B., A Generalized Compactifying Garbage Collector, Comput. J., 15, 3(Aug 1972), pp. 204-208

Weissman, C., Lisp 1.5 primer, Dickenson Publ., Belmont, Calif., 1967

Weizenbaum, J., Symmetric list processor, ACM 6, 9(Sept. 1963), pp. 524-544

Weizenbaum, J., Recovery Of Reentrant List Structures In

SLIP, CACM 12,7(July,1969) pp. 370-372

Winston P.H., Horn E.K.P. LISP, Addison-wesley, Mass, 1981

Wirth, Niklaus, Algorithms + data structures = Programs, Englewood Cliffs, N.J., Prentice Hall, 1976

Wise, D.S., Morris garbage compaction algorithm restores reference counts, ACM Trans. Programm. Lang. Syst. 1, 1(July 1979), pp. 115-120

Wodon, P.L., Methods of garbage collection for Algol 68, in Algol 68 implementation, J. E. L. Peck (Ed.), North-Holland, Amsterdam, 1971, pp. 245-262

Zave, D.A.. A fast compacting garbage collector, Inf. Process. Lett. 3, 6(July 1975), pp. 167-169

by resetting all counters to zero. The counters of the accessible cells are restored during the marking phase of the collection by incrementing a cell's counter every time the cell is visited. The collection reclaims inactive circular list structures and cells with maximum refcount which have become unreachable. The point of this hybrid scheme is to keep track of very small refcounts between necessary invocation of garbage collection so that nodes which are allocated and quickly abandoned can be returned to available space, delaying necessity for garbage collection.

The hybrid approach of garbage collection and reference counting suggested by Deutsch and Pobrow is particularly applicable to LISP [see DEUTSCH 76]. According to Clark [CLARK 77], it is based on statistical evidence that in most LISP programs, most refcounts (about 97 percent) are one. They propose three hash tables [POBROW 75]:

(1) The multiple reference table (MRT): Its key is a cell address and the associated value is the cell's refcount. Only cells whose refcounts are two or greater are listed in the MRT.

(2) the zero count table (ZCT) contains the addresses of cells whose refcount is zero. These cells may be of two types, those which are referred to only by the variables