

<OFFICE-2, CLEN, FILE-DOC.AUG;2,>, 30-Jan-81 09:12 CLEN :;;;

this updates <27292>

Introduction

1

The AUGMENT file system evolved from the NLS file system. This document updates the NLS file documentation in journal item (27292). The updates reflect the changes made to the NLS file system to create the AUGMENT file system.

1a

AUGMENT operates on a heirarchical, random file system with several unique features evolved over the years that make possible the efficient online interaction used by the OAD community. Having information stored within separate structure and data blocks aids in rapid movement within and between AUGMENT files; a "partial copy" locking mechanism provides security against attempted modification of a file by more than one user at the same time and provides a high degree of backup security against system failure or user error. This appendix includes a technical description of the file system as well as a discussion of motivating factors leading to its implementation. The design of the file system provides room for further extensions, some of which are also examined.

1b

Discussion of the heirarchical structure of AUGMENT files at a user level, as well as a description of the user commands that permit movement through the files, may be found in [1].

1c

This appendix is a revision of an earlier document which described the NLS file system as of January, 1976 and is current to August 1980. The January 1976 additions to the NLS file system, included property lists and inferior trees, which are currently used in the new graphics subsystem and offer great potential for the creation of new user entities.

1d

General Considerations Leading to the Current Design

2

The format and structure of AUGMENT files were determined by certain design considerations:

2a

It is desirable to have virtually no limit on the size of a file. This means it is not practical to have an entire file in core when viewing or editing it.

2a1

The time required for most operations on a file should be independent of the file length. That is, small operations on a large file should take roughly the same time as the same operations on a small file. The user and the system should not be penalized for large files.

2a2

In executing a single editing function, there may be a large number of structural operations.

2a3

A random file structure satisfies these considerations. Each file is divided into logical blocks that may be accessed in random order.

2b

An early version of the file system was implemented on the XDS-940. Minor changes in the logical structure of the file system were made in the conversion of the system from the XDS-940 to the PDP-10 for two reasons:

2c

1) The current OAD programming language, L10, is more powerful than the several languages it replaces, MOL and the SPLs. L10 permits special purpose constructions anywhere in its code. It is a higher level language and provides greater compiler optimization.

2c1

2) An effort has been made to further modularize the functions within the system to ease development by a team of programmers.

2c2

In Winter 1975 extensions to the file system were made introducing property lists as data elements at each structural node. The first use of this capability was in the recently developed graphics subsystem. Further discussion of these changes may be found below.

2d

Reliability and the AUGMENT File System

3

The reliability and security of file data both against system crashes and in face of the possibility of attempted simultaneous modification by more than one user were central goals in the design of the AUGMENT file system. An attempt was made to minimize the amount of work which would be lost due to both hardware and operating system difficulties.

3a

Unlike the sequential file systems of some editors which require copying large sections of a file whenever an edit is made, AUGMENT modifies copies of pages in which structural or data changes are made: all data in the original file is secure and a minimum of unaffected data is copied. Still other editors maintain recent changes in a dynamic buffer which may not be incorporated into the file in the event of a system crash; in AUGMENT, barring a major hardware collapse, all changes other than those specified by the command being processed are present in the copied pages. Again, the original file is untouched.

3b

Other techniques to assure high reliability have been used such as organizing the code and sequence of operations in a way to minimize time windows of high vulnerability.

3c

An important problem in an online team environment such as that at OAD involves group collaboration on the same data files. The current file system permits multiple readers and a single writer to a file. The person obtaining write access to a file locks it in a manner described below; no other user is then permitted to write on the file, though they may read the original material. Readers without write access do not see the changes of the user currently editing the file until the file is explicitly "updated," causing the incorporation of edits and the unlocking of the file.

Thus there can be no conflict between the edits of more than one writer.

3d

Details on the partial copy locking mechanism which implements these features of the AUGMENT file system are discussed below in section (XXX).

3d1

Recent Extensions to the AUGMENT File System

4

OAD recently extended the AUGMENT file system to include a list of data blocks (a property list) rather than the single textual data block which existed before. These property lists are now associated with AUGMENT structural nodes in the same manner that the single data block had been associated before. There is no restriction on the types of data nodes: for instance, graphic or numerical information may be possible as well combinations of data types within a single node. Additionally, data nodes may themselves have structure in the form of "inferior trees". The extended file system is upwardly compatible with the older file system: old files are still useable on the new file system without conversion.

4a

Short Technical Overview

5

This section gives a brief overview of the implementation of AUGMENT files. For more detail see section (XXX).

5a

Block Header and Types of Blocks

5b

An AUGMENT file is made up of a file header block, and up to a fixed number (currently 465) of 512-word (=equals one TENEX page) structure blocks (up to 95), and data blocks (up to 370).

5b1

There are several types of blocks, each with its own structure:

5b2

File header block--always page 0: contains general information about the file.

5b2a

Structure (ring) blocks--contain ring elements that implement the AUGMENT structure: there currently may be a maximum of 95 of these blocks, each containing 102 five-word ring elements. They may appear in file pages 6 through 100.

5b2b

Data blocks--contain the data (in linked property lists associated with structural nodes) of AUGMENT statements: each data block is composed of individual data elements made up of a five-word header followed by text strings or other data. There currently may be a maximum of 370 data blocks. They may appear in file pages 101 through 471.

5b2c

Miscellaneous blocks--not used in the current implementation.

5b2d

File Header Block

5c

In each file there is a header block that contains general information about that particular file. The header block remains in memory while the file is in use.

5c1

The file header is read into core by the procedure (nlbsesrc, ioexec, rdhdr). This procedure checks for the validity of certain keywords. If the file is locked and has a partial copy, the header is read in from the partial copy. If the partial copy header block is invalid in the key spots, the file is unlocked and the header read in from the original file. If that is bad, the file may be initialized.

5c2

RDHDR sets the value of the FILENO-th element in the table FILEHEAD. FILENO is the AUGMENT file number of the file. (It is an index into the file status table that provides, among other things, a correlation between JFNs for the original and partial copy and the single AUGMENT file number).

5c3

Procedures in (nlbsesrc, filmp,) are responsible for reading, manipulating, creating, garbage collecting, and storing into ring blocks and ring elements within those blocks, and data blocks and statement data blocks within them.

5c4

Structure Blocks -- Ring Elements

5d

Conceptually an AUGMENT file is a tree. Each node has a pointer to its first subnode and a pointer to its successor. If it has no subnode, the sub-pointer points to the node itself. If the node has no successor, the successor pointer points to the node's parent. (These conventions are used to aid in providing a set of primitives for rapidly moving around in AUGMENT files.) Each node is currently represented by a ring element. These ring elements point in turn to the first data block in the node's property list.

5d1

Structure blocks contain five-word ring elements with a free list connecting those not in use.

5d2

Data Block -- Property Lists and (Textual) Statement Data Blocks

5e

Data blocks are composed of variable sized elements called (Textual) Statement Data Blocks (SDBs) that contain the text of AUGMENT statements and other types of data elements. Other data element types are currently used in the AUGMENT graphics system though the number of available types and uses may be easily extended. All data elements have a five word header followed by data appropriate to the element type. Each SDB has this five-word header with node related information followed by the text made up of 7-bit ASCII characters packed five to a word. New data elements are allocated in the free space at the end of a data block page. Data elements no longer in use

(because of editing changes) are marked for garbage collection when the free space is exhausted.

5e1

Data elements associated with node are linked together in a property list. This property list may in turn have a structured inferior tree associated with it; the nodes on the inferior tree structure of a data element may also have associated property lists. This feature may prove to be useful in the creation of a comment entity in AUGMENT for comments associated with a particular AUGMENT statement.

5e2

Statement (or String) Identifiers (STIDS) and Text Pointers

5f

A statement identifier (STID) is a data structure used within AUGMENT to identify AUGMENT statements (structural nodes) or strings.

5f1

If the string is in an AUGMENT statement, the STID contains a file identifier field (STFILE) and a ring element identifier (STPSID).

5f1a

The presence of a file identifier within the STID permit all editing functions to be carried out between files.

5f1b

Text pointers are two-word data structures used with the string analysis and construction features of L10. They consist of an STID and a character count.

5f2

Locking Mechanism -- Partial Copies

5g

The AUGMENT file system under TENEX provides a locking mechanism that protects against inadvertent overwrite when several people are working on the same file. Once a user starts modifying a file, it is "locked" by him against changes by other users until he deems his changes consistent and complete and issues one of the commands: Update File, Update File Compact, or Delete Modifications, which unlock the file. A user can leave a file locked indefinitely--this protection is not limited to one console session.

5g1

When a file is locked (is being modified), the user who has modification rights sees all of the changes that he is making. However, others who read the file will see it in its original, unaltered state. If they try to modify it, they will be told that it is locked by a particular user. Thus the users can negotiate for modification rights to the file.

5g1a

This feature is implemented through the use of flags in the status table in the File Header and through the partial copy mechanism.

5g2

All modifications to a file are contained in a partial copy file. These include modified ring elements and data blocks. 5g2a

Any file page that is to be and that is not in the partial copy (discovered through a write pseudo-interrupt) is copied into the partial copy. All editing takes place there. The TENEX user-settable word in the FDB (TENEX file data block) for the original file contains locking information. 5g2b

The AUGMENT Update file command merely replaces those structure and data pages in the original file that have been superseded by those in the partial copy, unlocks the file, and deletes the partial copy. For Update file old, this is done in the original file; for Update to new version, the pages are mapped to a new file from the original or partial copy where necessary. The Update file compact command garbage collects unused space; the update file command does not. 5g3

Core Management of File Space 5h

When space is needed for more data, the following steps are taken, in order, until enough is found to satisfy the request (See (nlsbesrc, filmp, nwrngb), (nlsbesrc, filmp, newdb), and related routines): 5h1

1) Core-resident pages are checked for sufficient free space. 5h1a

2) Other pages are checked for free space. If one has sufficient space, it is brought in. 5h1b

3) If garbage collection on any page in the file will yield a page with sufficient free space, then the page that will give the most free space is brought into core and garbage-collected; otherwise a new page is created. 5h1c

Note on Fields in AUGMENT Records and Other L10 Language Features

6a

Several parts of this section are taken directly from record declarations in the code of the AUGMENT system written in the L10 programming language.

6a1

Record declarations in the L10 language serve as templates on data structures declared in the system. Byte pointer instructions are dropped out by the compiler permitting access to specified parts of the array. Multiword records are filled from the lowest to the highest address of the array. Within words, bits are allocated from the first bit on the right. If several fields fail to fill a 36-bit word and the next field definition would go over the remaining bits in the word, the field is allocated in the next word available.

6a2

Example:

6a2a

Bit 0 is the leftmost bit in the word; bit 35 the rightmost. Suppose there is a record declaration of the form:

6a2a1

```
(newrecord) RECORD % A two word record %
  field1[10], %bits 26 through 35 (rightmost) of
  first word%
  field2[25], %bits 1 through 25 of first word %
  field3[15]; %bits 21 through 35 of second word
  (field would not fit in remainder of first word%
DECLARE array[2];
```

6a2a1a

6a2a1a1

6a2a1a2

6a2a1a3

6a2a1b

There may be code within a program of the form:

6a2a2

```
variable _ array.field2;
array.field3 _ 20;
```

6a2a2a

6a2a2b

In L10, false is zero and true is nonzero.

6a3

See the L10 manual for further information.

6a4

Block Header and Types of Blocks

6b

An AUGMENT file is made up of a file header block page and up to a fixed number (currently 465) of 512-word (= one TENEX page) structure block pages (up to 95) and data block pages (up to 370).

6b1

Each page has a two-word header telling the type and giving the file page number and an index into a core status table. The record declaration from (nlsbesrc, brecords,) follows:

6b2

```

(fileblockheader) RECORD %fbhdl = 2 is length%           6b2a
  fbnull[36],      %unused%                               6b2a1
  fbnd[9],         %status table index%                   6b2a2
  fbpnum[9],       %page number in file of this block%   6b2a3
  fbtype[5]:       %type of this block (types declared in
  (nlsbesrc, bconst,))                                   6b2a4
    hdbtyp = 0     = header                               6b2a4a
    sdbtyp = 1     = data                                 6b2a4b
    rngtyp = 2     = ring                                 6b2a4c
    jnktyp = 3     = misc (such as keyword, viewchange, etc.
    Not currently used.)%                                 6b2a4d

```

PAGE HEADER BLOCK

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
X free                                                                 X
X 36                                                                 X
X-----X
X free          * Type * Page   * Status X
X               *      * Number * Table X
X 13           * 5   * 9     * 9     X
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

```

6b3

There are several types of block pages, each with its own structure.

6b4

File header block pages--always page 0: contains general information about the file.

6b4a

Structure (ring) block pages--contain ring elements that implement the AUGMENT structure: there currently may be a maximum of 95 of these blocks, each containing 102 five-word ring elements. They may appear in file pages 6 through 100.

6b4b

Data block pages--contain the data properties of AUGMENT statements: each data page has properties with five-word headers followed by data (text, graphics instructions, etc.). There currently may be a maximum of 370 data pages. They may appear in file pages 101 through 471.

6b4c

Miscellaneous blocks--not used in the current implementation.

6b4d

In each file, there is a header block page that contains general information about that particular file. The header block remains in memory while the file is in use.

6c1

FILE HEADER CONTENTS (taken from (nlsbesrc, bdata,)):

6c1a

```

DECLARE EXTERNAL                                     6c1a1
%...file header...%                                 6c1a2
% DONT CHANGE THE ITEMS IN THE HEADER %            6c1a2a
(filhed) EXTERNAL [4] ;                             6c1a2b
% these extra words may be taken for additions to
  header%                                           6c1a2b1
(fprot) EXTERNAL ; % protection word %             6c1a2c
(fcredt) EXTERNAL ; % file creation date %         6c1a2d
(nlsvwd) EXTERNAL = 1 ; % nls version word (keyword)
%                                                   6c1a2e
(sidcnt) EXTERNAL ; %count for generating SID*s%   6c1a2f
(finit) EXTERNAL ; % initials at last write %      6c1a2g
(funo) EXTERNAL ; % user number (file owner) %     6c1a2h
% if <0, RH is pointer to string in fileheader%    6c1a2h1
(lwtim) EXTERNAL ; % last write time %             6c1a2i
(namdl1) EXTERNAL ; % left name delimiter default
character %                                         6c1a2j
(namdl2) EXTERNAL ; % right name delimiter default
character %                                         6c1a2k
(rngl) EXTERNAL ; % upper bound on data ring file
blocks %                                           6c1a2l
(dtbl) EXTERNAL ; % upper bound on data file
blocks %                                           6c1a2m
(rfbs) EXTERNAL [6] ; % start of random file block
status tables %                                    6c1a2n
(rngst) EXTERNAL [95] ; % ring block status table % 6c1a2o
(dtbst) EXTERNAL [370] ; % data block status table % 6c1a2p
(mkrtxn) EXTERNAL = 20 ; % marker table maximum
length %                                           6c1a2q
(mkrtbl) EXTERNAL ; % number of markers in marker
table %                                             6c1a2r
%each marker takes MKRL words%                     6c1a2r1
(mkrtb) EXTERNAL [20] ; % marker table %           6c1a2s
(filhde) EXTERNAL ; %end of the file header%       6c1a2t

```

Notes on File Header

6c1b

The file header is read into core by the procedure (nlsbesrc, ioexec, rdhdr). This procedure checks for the validity of certain keywords. If the file is locked and has a partial copy (the file that includes current modifications--see below), the header is read in from the partial copy. If the partial copy header block is invalid in the key spots, the file is unlocked and the header read in from the original file. If that is bad,

the file may be initialized. RDHDR sets the value of filehead[fileno] where fileno is the AUGMENT file number of the file (an index into the file status table which provides, among other things, a correlation between JFNs for the original and partial copy and the single AUGMENT file number; see description of the file status table below.)

6c1b1

(nlsbesrc, ioexec, setfil) initializes a file header.

6c1b2

It should be noted that fields within a file header are accessed by full word indexing rather than by record pointers for speed. Thus we have the following typical code (from (nlsbesrc, filmnp, crepr2)) that reads the default name delimiters from an AUGMENT file header:

6c1b3

```
. 6c1b3a
. 6c1b3b
. 6c1b3c
% Must calculate the delimiters % 6c1b3d
IF stid.stpsid = origin THEN 6c1b3e
  BEGIN %use standard delimiters for that file% 6c1b3e1
  fhdlc _ filehead[stid.stfile] - $filhed; 6c1b3e2
  sdbnew.slndml _ [fhdlc + $namdl1]; 6c1b3e3
  sdbnew.srnmdl _ [fhdlc + $namdl2]; 6c1b3e4
  END 6c1b3e5
. 6c1b3f
. 6c1b3g
. 6c1b3h
```

Also, code from (nlsbesrc, ioexec, rdhdr) that gets the address of the word in core that contains the nls version word for the file whose header has been read in order to check its validity:

6c1b4

```
. 6c1b4a
. 6c1b4b
&vwd _ (header _ filhdr(fileno)) - $filhed + $nlsvwd; 6c1b4c
filehead[fileno] _ header; 6c1b4d
. 6c1b4e
. 6c1b4f
```

The file header is initialized by (nlsbesrc, ioexec, rdhdr) which fills up contiguous words declared in (nlsbesrc, bdata,) and then moves the contents of those words to page zero of the file.

6c1b5

FILE HEADER BLOCK (FULL WORDS)

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
X free[4] X X Max structure pages X
X-----X X-----X
X Protection word X X Max data pages X
X-----X X-----X
X Creation data X X Start of block tables[6] X
X-----X X-----X
X Version Number (=1) X X Ring block status table[95]X
X-----X X-----X
X SID Count X X Data blk status table[370]X
X-----X X-----X
X Initials last write X X Marker table size (=20) X
X-----X X-----X
X File Owner X X Marker table[20] X
X-----X X-----X
X Time last write X X End of file header X
X-----X X-----X
X Left name delimiter X X X
X-----X X-----X
X Right name delimiter X X X
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

```

6c1c

Procedures in (nlsbesrc, filmp,) are responsible for reading, manipulating, creating, garbage collecting, and storing into ring blocks and ring elements within those blocks, and data blocks and statement data blocks within them.

6c2

The random file block status tables appear in the file header. There is one word per ring block or data block page. Each entry contains the following: record declaration and comments from (nlsbesrc, brecords,).

6d1

(rfstr) RECORD % random file block status record %	6d1a
rfaxis [1], % true if the block exists in the file %	6d1a1
rfpart [1], % true if block comes from partial copy %	6d1a2
rfnull [2], % unused %	6d1a3
rfused [10], % used word count for the block %	6d1a4
rffree [10], % free pointer for the block %	6d1a5
rfcore [9]; % 0 then not in core, else page index %	6d1a6
% unallocated if entirely zero %	6d1a6a

BLOCK STATUS TABLE ENTRIES (STRUCTURE OR DATA)
XX
X * Page index else * Free * Used * * Part-*ExistX
Xfree* =0 if not in * pointer * word *free* ial * ? X
X * core * or * count * * copy?* X
X * * * count * * * * * X
X 3 * 9 * 10 * 10 * 2 * 1 * 1 X
XX

6d1b

Notes on Random File Block Status Tables

6d2

The table RFBS in the file header is broken into two sections, each of which contains a collection of records of the above type. The first section includes RNGM entries from RFBS[RNGBAS] up to and including RFBS[RNGBAS+RNGM-1] and contains information about the ring block pages in the file. (RNGBAS is currently 6 and is the first page in a file that may be a ring page; RNGM is currently 95 and is the maximum number of ring block pages permitted.)

6d2a

The second section includes DTBM entries from RFBS[DTBBAS] up to and including RFBS[DTBBAS+DTBM-1] and contains information about the data block pages in the file. (DTBBAS is currently 101 and is the first page in a file that may be a data block page; DTBM is currently 370 and is the maximum number of data block pages permitted.) The entry RFBS[RNGBAS+i] may also be referenced as RNGST[i]; likewise RFBS[DTBBAS+i] may be referenced as DTBST[i]. The index in RFBS is the actual page number of a data page in the file.

6d2b

A pointer to a data element or property (PSDB) consists of a nine-bit data page number in the range [0,DTBM) and a nine-bit displacement from the start of the page. The

variable DTBL is maintained in each file header as the current upper bound on allocated data pages for that file. This is used to limit the search for a location for a new data element. The variable DBLST contains the index of the block from which a property was last allocated or freed.

6d2c

A pointer to ring element (PSID) consists of a nine-bit ring page number in the range [0,RNGM) and a nine-bit displacement from the start of the page. The variable RNGL is maintained in each file header as the current upper bound on allocated ring pages for that file. This is used to limit the search for a location for a new ring block. The variable RNGST contains the index of the page from which a ring was last allocated or freed.

6d2d

These blocks contain five-word ring elements with a free list connecting those not in use.

6e1

(ring) RECORD % *** ringl is length% % from (compsrc, rt-main,)
%

rsub[18],	%psid of sub of this statment%	6e2
rsuc[18],	%psid of suc of this statement%	6e2a
rsdb[18],	%psdb of sdb for this statement%	6e2b
rinst1[7],	%DEX interpolation string-- scratch space%	6e2c
rinst2[7],	%DEX interpolation string-- scratch space%	6e2d
rdummy[1],	%DEX dummy flag-- scratch space%	6e2e
repet[3],	%DEX repetition-- scratch space%	6e2f
rhf[1],	%head flag, true if this is head of plex%	6e2g
rtf[1],	%tail flag, true if tail of plex%	6e2h
rnamef[1],	%name flag, true if statement has a name%	6e2i
rtorigin[1],	%inferior tree origin flag, true if origin%	6e2j
rnull[1],	%unused%	6e2k
rnameh[30],	%name hash for this statement%	6e2l
rsid[30],	%statement identifier%	6e2m
%although only need four words, use five so that have room to grow%		6e2n
		6e2o
		6e2p

RING ELEMENT

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
X PSID of Successor          * PSID of Substatement (Down)  X
X 18                          * 18                          X
X-----X
X Scratch space used by DEX * PSDB (pointer to data block) X
X 18                          * 18                          X
X-----X
Xfree* Name Hash              *free*org *name*tail*headX
X 1 * 30                      * 1 * 1 * 1 * 1 * 1 X
X-----X
X free * Statement Identifier X
X 6 * 30                      X
X-----X
X free X
X 36 X
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
    
```

6e3

PSIDs and PSDBs are pointers to other ring or data blocks in a file. They have two nine-bit fields: one (stblk) is a page index; the other (stwc) is a word displacement within that page. Procedures in (nlsbesrc, filmnp,) permit the traversal of a file's structure.

6e4

Given an STID (see below), one may use the primitive procedures in (nlsbesrc, filmnp,)--e.g., (nlsbesrc, filmnp, getsuc)--or the more elaborate procedures in that file--e.g., (nlsbesrc,

filmp, getnxt)--to move around to related ring elements and retrieve or change (display or edit) relevant data. 6e5

There are two "fixed" values for PSIDs for special statements: 6e5a

The PSID of the origin statement is always 2. 6e5a1

The entire STID (and hence PSID) of the end of a file is endfil (= -1), which does not correspond to any real statement in the file, but which is returned by the "get" procedures in filmp to indicate the end has been reached or an error has been found. 6e5a2

Some other conventions implemented in the file structure make possible special features in AUGMENT: 6e5b

The successor of a statement with no real successor is its "parent." 6e5b1

The substatement of a statement with no sub is itself. 6e5b2

The origin is at a unique level; thus statement 1 is the sub of the origin. 6e5b3

Property lists are made up of linked lists of property data blocks. An example of a property is the Statement Data Block (SDB) which contains the text of an AUGMENT statement.

6f1

Each property has a five-word general header with the following information. There then follows data appropriate to the particular property type. For example, (Textual) Statement Data Blocks (SDBs) contain the text in AUGMENT statements; this text follows the property header and is composed of seven-bit ASCII characters packed five to a word, in a variable length block. New properties are allocated in the free space at the end of a data page. Properties no longer in use (because of editing changes) are marked for garbage collection when the free space is exhausted.

6f2

```
(sdbhead) RECORD %sdbhdl is length% % from (compsrc, rt-main,)
```

%

```

sgarb[1],      %true if this sdb is garbage%           6f3
length[9],     %number of words in this sdb%           6f3a
schars[11],    %number of characters in this statement% 6f3b
slnmdl[7],     %left name delimiter for statement%     6f3c
srnmdl[7],     %right name delimiter for statement%    6f3d
spsid[18],     %psid of the statement for this sdb%    6f3e
sname[11],     %position of character after name%      6f3f
stime[36],     %date and time when this sdb created%  6f3g
sinit[21],     %initials of user who created this sdb% 6f3h
sptype[15],    %property type of this data block%     6f3i
%                                                       6f3j
txttyp        = text data block (SDB)                 6f3j1
dhdtyp        = diagram header block                  6f3j2
segtyp        = segment data block                    6f3j3
%                                                       6f3j4
%                                                       6f3j5
spsdb[18],     %PSDB of the next property data block 0=tail% 6f3k
sitpsid[18];   %PSID to head of inferior tree if any%  6f3l
%sgarb and length must be in the first word of the header 6f3m
for newsdb%    6f3n

```

DATA BLOCK HEADER

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
X      *Right name *Left name  *Character*Block *Garb-X
X free  *delimiter *delimiter  *count   *length*age? X
X 1     * 7        * 7         * 11    * 9     * 1    X
X-----X
X free * Position of char * PSID pointer to ring element X
X 7    * 11  after name   * 18   for this statement   X
X-----X
X Creation time                                     X
X 36                                             X
X-----X
X Property      * Authors initials                 X
X 15  type      * 21                               X
X-----X
X PSID of inferior tree * PSDB of the next property X
X 18                                             * 18   X
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

```

6f4

STATEMENT DATA BLOCK (SDB*S) Text type block

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
X Data block header                                     X
X 5 full words                                         X
X-----X
X Text                                                 X
X Block length - 5 words of 5 characters each         X
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

```

6f5

Statement (or String) Identifiers (STIDS) and Text Pointers 69

A statement identifier (STID) is a data structure used within AUGMENT to identify AUGMENT statements (structural nodes) or strings. 691

If the string is in an AUGMENT statement, the STID contains a file identifier field (STFILE) and a ring element identifier (STPSID). (See PSID description above under ring elements.) 691a

The presence of a file identifier within the STID permit all editing functions to be carried out between files. 691b

Procedures in (nlsbesrc, filmp,) permit traversal through the ring structure of a file given an STID. See, for example, (nlsbesrc, filmp, getsuc), which gets the STID of the successor of a statement; see also (nls, filmp, getsdb), which returns the STDB for the statement whose STID is provided as an argument. (An STDB has, like an STID, a file number field and a pointer to the textual property block in the property list, a STPSDB). Additional primitives are available for other data properties. 692

Text pointers are used with the string analysis and construction features of L10. They consist of an STID and a character count. 693

Other Relevant Arrays

6h

The following arrays are used in system core and file management. They are described here to facilitate the study of the AUGMENT file-handling code.

6h1

Filehead

6h2

An array of pointers (each contained in a single word) to the file headers of files currently in use is FILEHEAD. At present, up to 26 files (and their partial copies, if any) may be open simultaneously.

6h2a

CORPST (Core Page Status Table) and CRPGAD (Core Page Address Table)

6h3

The array CORPST provides the correspondence between the 100 (octal) pages in core reserved for file pages and user program buffer and the pages in files that are currently loaded into core. (This is really a maximum of 100 octal since the user program buffer may be enlarged into this area; the maximum is given by $RFP_{MAX} - RFP_{MIN} + 1$.)

6h3a

(corpgr) RECORD %core page status record%	6h3a1
ctfull [1], % true if the page is in use %	6h3a1a
ctfile [5], % file to which the page belongs %	6h3a1b
ctpnun [9], % page number within the file %	6h3a1c
ctfroz [3]; % number of reasons why frozen %	6h3a1d
% a single word; gives status for a given core page for random files %	6h3a1d1

The array CORPST is the core page status table and is made up of instances of the above record. $(\max RFP - \min RFP + 1)$ gives the number of core pages that may contain file pages. The core pages are located at positions indicated by the array CRPGAD (core page address). CORPST is indexed by numbers in the range (RFP_{MIN}, RFP_{MAX}) . The elements in this array are actual addresses. The starting location of page k is given by $crpgad[k]$. RFP_{MIN} is initialized to be 7; six pages are initially allocated for a user program buffer. See (sssrc, programs,) for the procedure that changes these limits.

6h3b

FILST (File Status Table)

6h4

An AUGMENT file number provides an index into the FILST, the file status table. This 100-word array is made up of 25 four-word entries and contains the following information for files of interest that have AUGMENT file numbers at any time (these may or may not at that time be open; they do, however, have JFNs.) The information comes from the record declaration in (nlbsesrc, brecords,):

6h4a

(filstr) RECORD %File status table record. entry length			
= filstl = 4, max no. entries = filmax = 25%			6h4a1
flexis [1], % true: entry represents an existant			
file %			6h4a1a
flhead [9], % crgpad index of the file header %			6h4a1b
flbrws [1], % this file in browse mode %			6h4a1c
fllock [1], % file was locked by another user			
when loaded %			6h4a1d
flpcread [1], % PC read only-- write open failed			
(openpc) %			6h4a1e
flaccm [8], % file access mask %			6h4a1f
fldirno [12], % directory number for the original			
file %			6h4a1g
flnoclos [1], % do not close this file %			6h4a1h
flinclude [1], % do not close this file: it is			
included %			6h4a1i
flhelp [1], % do not close this file: included			
by Help %			6h4a1j
flpart [18], % JFN for the partial copy %			6h4a1k
flbpart [18], % JFN for the browse partial copy %			6h4a1l
florig [18], % JFN for the original file %			6h4a1m
flastr [18], % address of the file name string %			6h4a1n
flpcst [18], % address of partial copy name			
string %			6h4a1o
flbpcst [18], % address of browse partial copy			
name string %			6h4a1p
flnsw [18]; % address of nsw filename string %			
			6h4a1q
(filstl) EXTERNAL = filstr.SIZE;			6h4a1q1
(filmax) EXTERNAL = 25;			6h4a1q2

Primitives for Use with Basic AUGMENT File Entities	6i
Introduction	6i1
The following primitives will be available for manipulation of basic file entities. While they make use of even more basic procedures, most programmers should have no reason for accessing lower level routines. These primitives and lower level procedures live in the file FILMNP.	6i1a
Property types must be assigned numbers by DAD. Code for management and portrayal of properties not generally available or useful for all AUGMENT users will be managed and written by the prime users. The procedures listed below will provide access to property blocks and nodes in the files.	6i1b
The code which manages graphics file entities lives, currently, in the graphics subsystem.	6i1b1
Entity types	6i2
Primitives will be available to operate on the following file entity types:	6i2a
NODE -- a ring element and its associated data contained in a property list.	6i2a1
PROPERTY -- a data block and any associated inferior tree within the property list associated with a node.	6i2a2
INFERIOR TREE -- structure and data associated with a property block.	6i2a3
An example of the use of an inferior tree may be found in the graphics subsystem in which diagrams have structure reflected by the existence of this inferior tree. Another possible use could be for imposing the structure (AUGMENT Plex-like in nature) of comments associated with a statement's text. Normal AUGMENT structural procedures for examining structure and modifying it at the file level may be used at the inferior tree level as well.	6i2a3a
Note that while no direct primitives are provided for operating on property lists or portions of them, such primitives exist at lower levels. It is not felt that higher level primitives for such entities are necessary. The operations listed below follow the currently existing examples for text nodes in AUGMENT files.	6i2b
Operands and procedures	6i3

READ -- Most read functions are dependent on the property type and are to be managed by formatters and other specific application code. Thus a set of "get" and "set" routines are available for examining and setting fields in the statement text nodes and similar procedures exist in the graphics subsystem. A general primitive to load particular property types into core is provided. Also, the usual procedures for moving around in structure will be available. 6i3a

lodprop (stid, proptype) -- 6i3a1

Loads the indicated property block into core. Returns three items: first is FALSE if error, page number in core if success; second is address of block in core (which must be frozen if you want to do anything with it!); third is stdb of property block 6i3a1a

Note change: as originally written, lodprop also took "occurrence" of property in list. we now will not permit more than one property of a particular type in a particular list. Multiple occurrences may be handled by a structural inferior tree hanging off the property block. 6i3a2

CREATE -- Allocate space for entity and link the blocks into existing structure and/or data. 6i3b

crenod (stid, rlevcnt) -- 6i3b1

Gets a new ring element with no associated data blocks and links it into the structure at the location specified by stid and rlevcnt (a relative level count: < 0 is down, =0 is successor, > 0 is up by rlevcnt levels). Returns stid of new ring or 0 if error. 6i3b1a

creprop (stid, proptype, length, data) -- 6i3b2

Builds a data block of property type proptype which must be a valid type assigned (and declared) by OAD and links it into the property list associated with the stid in the proper order (determined in the procedure linkprop). If such a property already exists in the node, we have an error: it must first be deleted. Returns stdb of new block or 0 if error. length is the length of the data and data is a pointer to an array of length words in which the data is stored. 6i3b2a

creit (stid, proptype) -- 6i3b3

Creates the origin of an inferior tree and links it to the data block property specified by stid and proptype. Returns 0 if error or stid of origin of inferior tree. 6i3b3a

DELETE -- Unlink entity from other structure and data.
 Release space. 6i3c

delprop (stid, proptyp) -- 6i3c1

deletes the property block and any associated inferior tree structure for the block proptype block of the indicated node. Returns TRUE if successful, 0 if not. 6i3c1a

delit (stid, proptype) -- 6i3c2

deletes the inferior tree of the indicated property block. Unlinks it and releases space. Returns True if successful, 0 if not. 6i3c2a

Currently no primitive exists to directly delete a node, though the primitives remgrp and delgrp perform this function together. The x-routines which implement structural deletes call these file system primitives. 6i3c3

MOVE -- Unlink entity and relink it at new location in file. 6i3d

movprop (stid, proptyp, destid) -- 6i3d1

Moves the property indicated from node specified by stid to node specified by destid. Accomplishes this by unlinking and relinking the block. If a property type of the type being moved exists at the destination, we have an error. Returns true if OK, 0 if error. 6i3d1a

movit (stid, proptype, destid) -- 6i3d2

moves the inferior tree associated with property block indicated by stid and proptype to the property block proptype associated with node destid. Returns true if OK, 0 if error 6i3d2a

X-routines currently exist to move and move filtered other structural entities. 6i3d3

COPY -- Allocate space for new entity, copy old entity, and Link the new entity into the file. 6i3e

copprop (stid, proptype, destid) -- 6i3e1

copies property block (and associated inferior tree if any) from block indicated by stid and proptype to a new block to be created on destid. Returns TRUE if OK. 0 if error 6i3e1a

citree (stid, proptyp, destid) -- 6i3e2

copies inferior tree of property block at node indicated by stid and proptype to the proptype block of destid. Returns TRUE if successful, 0 if error 6i3e2a

Primitives exist to do structural copies both in filtered and unfiltered modes. 6i3e3

REPLACE -- In keeping with the mode which exists for text statements, a replace primitive will not be provided for the inferior tree entity or the node entity. These functions may be accomplished using existing x-routines or primitives which delete nodes followed by a copy or create. 6i3f

reprop (stid, proptype, length, data) -- 6i3f1

replaces the property block indicated by stid and proptyp 6i3f1a

with a block with data as indicated. If length is the same as 6i3f1b

the length of data in the existing property block, a short cut may 6i3f1c

be taken and the data overwrites the old data. If, however, 6i3f1d

the length is different, a new block is built and linked in. 6i3f1e

The inferior tree is not replaced in any case: it remains the 6i3f1f

same. The inferior tree's pointer to the "owning" property 6i3f1g

block is changed to point to the new block. Uses filesc if this is a text block. 6i3f1h

References

7

(17b2) Douglas C. Engelbart and Staff of ARC. Computer-Augmented Management-System Research and Development of Augmentation Facility [Final Report]. Augmentation Research Center, Stanford Research Institute, Menlo Park, California 94025. APR-70. (5139.)

7a