

Reorganize

Start

Data structures:

Empty: SEQUENCE slot: [0..n] OF {empty, valid, badHash, dibbiesEmpty, dibbiesBadHash};

resolveList, dupList: LIST OF [currentPageNo, desiredPageNo.];

Read thru the file, counting the number of valid entries.

Should the size of the file change?

smaller

bigger

we expect this to be very rare, so just start at what will be the new eof and chuck valid entries over our shoulder into empty slots, in the order in which we find the empty slots.

extend the file, setting the new entries to "empty".

Build a sequence Empty the size of the new file and set all entries to "empty". Set the resolveList to NIL.

Go to the beginning of the file.

Is there another record? Yes.

is it empty? Yes.

No.

No.

Is it deleted? Yes. Write an empty record in the file.

No.

Must be valid. Does it have the right hash for this slot? Yes. Set the entry in the Empty seq to valid.

No. Make an entry in the resolve list. Set the entry in the Empty seq to badHash.

TURN ALL DELETED RECORDS INTO EMPTY RECORDS. ALL PAGES IN THE WRONG SLOTS GET NOTED ON THE RESOLVE LIST. SET UP EMPTY ARRAY WITH THE VALUES EMPTY, VALID, BADHASH AS APPROPRIATE.

Set the dupList to NIL.

Start at the beginning of the resolveList.

Is there another entry in the resolve list? Yes.

Is the desired slot in the Empty seq = empty or badHash?

No.

No, valid or dibbies. Yes.

Move our entry from the resolveList to the dupList. Set that slot to dibbiesEmpty or dibbiesBadHash.

SEE IF THE SLOTS THE PAGES ON THE RESOLVE LIST WANT ARE AVAILABLE. FIRST PAGE TO REQUEST A SLOT THAT IS CURRENTLY ANYTHING BUT VALID GETS DIBBIES ON THAT SLOT. LOSERS GO ON DUPLIST.

Start at the beginning of the resolveList.

Is there another entry in the resolve list? Yes.

Does the slot in Empty for the desired page = dibbiesEmpty? Yes.

No.

No, dibbiesBadHash.

Exch the two disk pages and update the other page's curPageNo in its entry in the dupList or resolveList. Write our page there. Remove our entry from the resolveList. Set slot to valid.

WRITE EVERYBODY ON THE RESOLVE LIST TO THE DISK. OVERWRITE EMPTY PAGES, EXCH WRONG HASH PAGES.

Start at the beginning of the dupList.

Is there another entry in the dupList? Yes. Calc the next best desired slot. Is it valid? Yes.

No.

No.

Is the desired slot = empty? Yes. Write our page there.

No. Exch the two disk pages and update the other page's curPageNo in its entry in the dupList. Remove our entry from the dupList. Set slot to valid.

FOR EACH ENTRY ON THE DUPLIST, CALC THE DESIRED SLOT, THEN SEARCH FORWARD (MAY WRAP TO BEGIN OF FILE) UNTIL FIND AN AVAILABLE SLOT.