HASHING SCHEMES FOR EXTENDIBLE ARRAYS

Extended Abstract

Arnold L. Rosenberg Larry J. Stockmeyer Mathematical Sciences Department IBM Watson Research Center Yorktown Heights, New York

ABSTRACT: The use of hashing schemes for storing extendible arrays is investigated. It is shown that extendible hashing schemes whose worst-case access behavior is close to optimal must utilize storage inefficiently; conversely, hashing schemes that utilize storage too conservatively are inevitably poor in expected access time. If requirements on the utilization of storage are relaxed slightly, then one can find rather efficient extendible hashing schemes. Specifically, for any dimensionality of arrays, one can find extendible hashing schemes which at once utilize storage well [fewer than 2p storage locations need be set aside for storing arrays having p or fewer positions] and enjoy good access characteristics [expected access time is 0(1), and worst-case access time is 0(log log p) for p- or fewer-position arrays]. Moreover, at the cost of only an additive increase in access time, storage demands can be decreased to $(1+\varepsilon)p$ locations for arbitrary $\varepsilon>0$. In fact, if one will abide a more drastic degradation of access efficiency, one can lower storage demands to p+o(p) locations.

1. INTRODUCTION

Conventional schemes for storing arrays are not readily extendible. For instance, in two dimensions, the familiar store-by-row scheme admits easy adjunction of new rows but only cumbersome appendage of columns. Such asymmetry in extendibility is not inevitable: there are computed-access schemes for storing arrays which are readily extended in any direction [1]. (An array storage scheme uses computed access if it assigns an address to an array position as a displacement from the address assigned to position <1,...,1>, the displacement being computed from the position's coordinates.) However, extendibility in array realizations does not come without cost. The most definitive illustration of the cost of extendibility is the study in [2] of the efficiency of storage utilization by extendible array realizations. It is shown in that paper that every d-dimensional extendible array realization must, for every integer p, "spread" some array having p or fewer positions over at least $0(p \cdot (\log p)^{d-1})$ storage locations. This bound suggests that the costs of unbridled extendibility are prohibitive; and it indicates that ways of overcoming the bound should be sought. Two avenues to a smaller lower bound are available: one can abandon the demand for unbridled extendibility, or one can abandon the restriction to computedaccess realizations. With regard to the former alternative, it is shown in [2] that, by restricting the patterns of expansions of arrays and by focusing only on arrays which conform to the restrictions, one can improve storage "spread" to O(p), irrespective of the dimensionality of one's arrays. With respect to the abandonment of computed access, two alternatives arise. If one's arrayprocessing algorithms are predominantly traversaloriented, then linked realizations of arrays ("orthogonal lists" in Section 2.2.6 of [3]) are an attractive alternative to computed-access realizations. Such realizations are at once easily extended and conservative of storage; only cp locations are needed for a p-position array, where c is a small integer. If, on the other hand, one's algorithms tend to access arrays by successive independent probes, then hashing represents a viable alternative to computed access, provided that hashing schemes can be found which are readily extended, conservative of storage, and not overly expensive to access. It is our purpose in this note to study the cost of extendibility in hashing schemes for extendible arrays.

Summary of Results. In Section 2, we describe how hashing schemes can be used to store extendible arrays; we use a simple bucket model of hashing schemes, with buckets organized as balanced search trees. We also define there the measures of efficiency that are the subject of our study. Section 3 is devoted to discussing a fundamental tradeoff. For any d-dimensional hashing scheme, the product of total space requirements and maximum bucket size for p-position arrays must grow at least as $p \cdot (\log p)^{d-1}$; this generalizes the result in [2] about computed-access realizations (which can be viewed as hashing schemes whose buckets never contain more than one array position). Various extreme cases are considered in Section 4. Hashing schemes whose worst-case access time does not grow with the size of the array being stored (computed-access realizations are included here) must utilize storage very poorly: total storage needed for p- or fewer-position d-dimensional arrays grows as $p \cdot (\log p)^{d-1}$. Conversely, hashing schemes that do not use successively more buckets as the arrays being stored grow suffer very inefficient access: expected access time for p- or fewerposition arrays grows no slower than log p. Surprisingly, similarly inefficient access plagues

hashing schemes that do use ever more buckets but do so in a very conservative way: included here are gap-free schemes which insist that, whenever bucket b>1 is used in storing array A, so also must be bucket b-1; here also are those schemes whose total storage demand for p- or fewer-position arrays is p+0(1) locations. In Section 5 we relax the extreme restrictions on storage allocation that plagued the schemes of Section 4, and we seek schemes that are good in both utilization of storage and time of access: we study hashing schemes that use O(p) buckets for p-position arrays. We show that any such linear hashing scheme must have worst-case access time of at least 0(log log p) for p- or fewer-position arrays (a dramatic improvement, if attainable, over Section 4's lower bound of log p for expected access time). The main message of Section 5, and, indeed, of the paper, is that this lower bound is achievable along with constant expected access time! Specifically, for every dimensionality d there is a d-dimensional extendible linear array-hashing scheme which, when storing arrays with at most p positions, has total storage demands of fewer than 2p locations,* has worst-case access time 0(log log p), and has expected access time c(d) ($c(2) \le 3$). The exact multiple of p in the expression for total storage demands is of little import here since, at the cost of only an additive increase in (both expected and worst-case) access time, this multiple can be brought below (1+ ε) for arbitrarily small ε >0. Indeed one can bring total storage demands down to p+o(p) if one is willing to suffer a more drastic degradation of efficiency of access (but no worse than $0(\log \log p))$.

<u>Related Work</u>. Knuth [3, Section 2.2.6] discusses both <u>sequential</u> (= computed-access) and <u>linked</u> allocation schemes for arrays; and he provides a comprehensive list of references to earlier work on such schemes. While scatter storage (= hashing) for arrays was undoubtedly considered and maybe even implemented in the 60's, we know of no reference predating [4] that discusses such schemes for storing arrays. The use of computed-access schemes for storing extendible arrays was considered first in [1]. The cost of such extendibility in terms of efficiency of storage utilization is studied in [2,5].

2. EXTENDIBLE ARRAY-HASHING SCHEMES

A. Notation

Let N denote the positive integers; and, for each neN, let N_n denote the set N_n = {1,...,n}. For arbitrary deN, N^d is the set of d-tuples of positive integers; we ambiguously let $\varepsilon = <1, \dots, 1>$, relying on context to specify the dimensionality of any instance of ε . For $\pi \in \mathbb{N}^d$, π_i (i = 1,...,d) is the ith coordinate of π ; thus, $\varepsilon_i = 1$ for all i. Finally, for any integer tuple π , $\Sigma(\pi) = \sum_i \pi_i$ and $\Pi(\pi) = \Pi_i \pi_i$.

B. Array Schemes and Hashing Schemes

In consonance with conventions for computedaccess array realizations [1,3] and for the arrayhashing functions of [4], we do not allow our hashing schemes to be data dependent; that is, we restrict attention to hashing functions that assign array positions to buckets using only a position's coordinates (and not its contents) to determine its bucket assignment. Accordingly, we can use as our notion of array the simple array schemes of [2,5] rather than any more elaborate representation of arrays.

- (2.1) The d-dimensional array scheme (array, for
 - short) of size $\langle n_1, \dots, n_d \rangle$ $(d, n_1, \dots, n_d \in \mathbb{N})$ is the set $A = \sum_{\substack{n \\ n_1 \\ n_d}} \langle n_d \rangle$. Each $\pi \in A$ is called a position of A.

Graphically, we envisage an array scheme as being imbedded in the positive orthant of the appropriate dimensional space, with its positions laid on the lattice points. Extensions to arrays (adjoining new rows and/or columns, in two dimensions) can be viewed as extensions to the discrete rectangular solid formed by the imbedding. (Note that the "rectangularity" resides in A's being the cross-product of the coordinate sets N_{n_4} .)

We view our hashing schemes as operating in the following simple manner. There is a bucket function b that assigns each position of the array A being stored to a bucket; for simplicity, we view the buckets as being addressed by positive integers, so b is a function from A into N. Of course several array positions will usually reside in the same bucket (or else b would be a computed-access realization as in [1,2]); the collisions caused by b's being many-to-one are resolved by organizing all the residents of each bucket into a balanced search tree [7, Section 6.2.3] by sorting them according to some key such as the lexicographic ordering of their coordinates. (By so exploiting the order inherent in the keys, one can appreciably reduce both expected and worstcase access time in external [= bucket-employing] hashing schemes. It is not obvious that internal hashing schemes can benefit from such exploitation, but Amble and Knuth [8] show that internal schemes can be so improved.) Thus we imagine the storage available to our hashing scheme to be divided into a collection of bucket addresses and a pool of chainable storage for the trees. See Figure 1.

The specific problem we address in this paper is to gauge the cost of extendibility in arrayhashing schemes. We lead up to our formal setting by considering the following scenario. We want to devise a bucket function b that will store any array of a given dimensionality d. We intend to use it as follows. We construct a system (say, a compiler or operating system) that includes an array-storage allocator based on b. As a customer (say, a program or a user) declares that his variable X will range over d-dimensional arrays having p or fewer positions, the system sets aside for X a block of m+n locations, call them $\{1, \dots, m, m+1, \dots, m+n\};$ here $m = \max\{b(\pi) \mid \pi \text{ is a }$ position of an array having p or fewer positions}, and n is the largest number of locations the system will ever need for constructing trees.

^{*} Each location is assumed capable of holding one datum, two pointers (to other locations in the same bucket), and a key used for searching.



The Array Scheme $N_4 \times N_4$ Stored by the Hashing Scheme b(i,j) = i+j-1. Starred position (= roots of trees) reside in "bucket" storage; the others reside in chainable storage

(Precisely, $n = \max\{q - \#b(A) | A \text{ is a } (q \le p) - position \}$ array}.) One can now store by hashing any array in X's domain, using the storage set aside by the system and using b as the bucket function. The "extend-ibility" in the described situation resides in the facts that (1) the same function b is used as the bucket function irrespective of what integer p appears in the declaration binding X; and (2) the scheme generated by the allocator will properly store any array in the domain of the variable X. The described scheme exhibits two types of stability: first, the computation required to store and/or access array positions does not change as the array being processed is extended; second, bucket assignments never change because of changes to the stored array. In fact, even if the declared value of p is changed dynamically, relatively minor rechaining within buckets can restore a stable configuration (especially if the pool of chainable storage is accessed from the high-address side). It is such stability under alterations to the stored array that we equate with "extendibility" of the hashing scheme. An alternative to our scenario might depict a different function b being used for each declared bound p. It is hard to see how any such system would exhibit the uniformity of scheme-generation or of extension and the algorithmic simplicity enjoyed by our system. The major question raised by this alternative is how much our uniformity costs. Theorems 5.3 and 5.4, our main results, suggest that the schemes generated by our imaginary system can attain a practical level of efficiency in both time of access and utilization of storage.

For simplicity, we henceforth identify a hashing scheme with its bucket function.

(2.2) A d-dimensional extendible (array-) hashing scheme is a total function $b: \mathbb{N}^d \to \mathbb{N}$ such that $b(\varepsilon) = 1$.

The normalizing assumption " $\flat(\varepsilon) = 1$ " simplifies certain computations in what follows but is not indispensable. (Cf. (2.4(a)) in the absence of this assumption.) Henceforth we shall study only (d>2)-dimensional extendible hashing schemes.

C. Measures of Efficiency

The efficiency of any hashing scheme is measured in terms of the scheme's demands on two resources, time and storage. However, deciding on what to measure is only half the battle; we must decide on how to estimate the consumption of these resources.

Access Time. We view the cost of computing a bucket function b as being negligible compared to the cost of searching even a modest size tree. (To make this assumption reasonable, we have tried to use only very simple bucket functions when establishing upper bounds.) Accordingly, we gauge the cost of accessing position π of array scheme A when A is stored by b as being the log of the population of π 's bucket; precisely,

$$\underline{\operatorname{Access}}_{p}(\pi; A) = \left\lceil \log(\#(\operatorname{Anb}^{-1}(\mathfrak{b}(\pi))) + 1) \right\rceil.$$

The justification for the logarithmic cost is that each bucket is organized into a balanced search tree. Positions added to an array via extensions are to be stored at the time of their first access and so stay within our logarithmic cost. The base of the logarithm used is immaterial in almost all of our results because of their "big-Oh" assertions. In the few cases (e.g., Theorems 5.3 and 5.4) where constants are crucial to the development, we use the base 2 logarithm appropriately dilated. (See Theorem A of [7, Section 6.2.3].)

Now that we have a way of gauging an extendible hashing scheme's accessing characteristics on a single array, we can formulate the first two functions of interest, which reflect a scheme's access behavior on arrays having no more than p positions. Focus on an extendible hashing scheme b.

- (2.3)(a) Worst-Case Access Time
 - a(p;b) = (the worst-case time needed by
 b to access a position of an
 array having at most p positions)
 = max{Access (π;A) | A has at most
 p positions, and πεA}.
 (b) Worst Expected Access Time
 - $\overline{\alpha}(p; \underline{b}) = (\text{the average time needed by } \underline{b} \text{ to}$ access a position of that p- or fewer-position array having worst $expected \text{ access time under } \underline{b})$ $= \underline{max}\{(1/q) \sum_{\pi \in A} \underline{Access}_{b}(\pi, A)\}$

A has $q \le p$ positions}.

Storage Requirements. Our view of extendible hashing schemes demands that we segregate the storage used for chaining (for building the trees that make up the bucket interiors) from the storage used for bucket headers (= the roots of those trees). If we did not partition storage in this way, a series of expansions of a stored array might demand for use as the header of bucket k a location already being used in the tree of bucket $l \neq k$. Such partitioning also is consonant with the scenario motivating our view of extendible hashing schemes. Finally, this division of storage makes the analysis of proposed hashing schemes or classes of hashing schemes immeasurably simpler than it would be if we used some more sophisticated method of collision resolution (e.g., linear probing or secondary clustering [8]). We formalize this division of storage by our method of measuring efficiency of storage utilization. Again, let b be an extendible hashing scheme.

- (2.4) (a) Bucket Storage Requirements
 - β(p;b) = (the number of storage locations over which b "spreads" the bucket headers [roots of search trees] when storing arrays having at most p positions)
 - $= \max\{\mathfrak{b}(\pi) \mid \Pi(\pi) \leq p\}.$

Note: π is a position of a p- or fewerposition array iff $\Pi(\pi) \leq p$.

- (b) Chainable Storage Requirements
 - σ(p;b) = (the amount of storage that is needed to store the bucket interiors [the nonroot nodes of the search trees] when storing arrays having at most p positions)
 - = $\max\{q-\#b(A) | A \text{ has } q \le p \text{ positions} \}$.
- (c) Total Storage Requirements
 - τ(p;b) = (the total amount of storage that must be set aside to store arrays having no more than p positions)
 - $= \beta(\mathbf{p}; \mathbf{b}) + \sigma(\mathbf{p}; \mathbf{b}).$

The aim of this paper is to discover relationships among the five efficiency measures of (2.3) and (2.4). In stating results, we shall fix the hashing scheme b by either quantification or instantiation; we shall then be free to consider our five functions as functions of p alone, as in "For all hashing schemes b, $\sigma(p;b) = O(p)$," or

" $\overline{\alpha}(p; b_0) = 0(1)$." [As an aside, the former assertion is obviously true; and we shall present a scheme in

Section 5 for which the latter assertion is true.]

The following lemma will be used in our investigation.

- (2.5) For all dimensionalities $d \in \mathbb{N}$ and $p \in \mathbb{N}$, $S_{d}(p) = \{\pi \in \mathbb{N}^{d} | \Pi(\pi) \leq p\}.$

3. A FUNDAMENTAL TRADEOFF

As a vehicle for introducing the reader to extendible hashing schemes, and as a prologue to a challenging research problem, we present a fundamental tradeoff in storage use by extendible hashing schemes. The tradeoff involves the total storage function $\tau(p;b)$ and the function now defined.

We prepare the reader for the tradeoff statement via some examples.

- (3.2) <u>A Computed-Access Realization</u>. If $b: \mathbb{N}^{d} \to \mathbb{N}$ is one-to-one, then b is an extendible array realization in the sense of [1,2,5]. It follows from Lemma 2.1, then, that $p \cdot (\log p)^{d-1} = O(\tau(p;b))$; and it is tautologous that <u>Size(p;b)</u> $\equiv 1$.
- (3.3) <u>A Linked-Allocation Scheme</u>. If $b(\pi) \equiv 1$, then b is a pure chaining scheme. It is immediate, then, that $\tau(p;b) = \underline{Size}(p;b) = p$.
- (3.4) Define $b: \mathbb{N}^{d} \to \mathbb{N}$ by $b(\pi) = \Sigma(\pi) d+1$; the two-dimensional (d=2) version of b is illustrated in Figure 1. It is not difficult to verify that $\beta(p;b) = p$, and $\sigma(p;b) =$ $p-dp^{1/d}+d-1$, so that $\tau(p;b) = 2p-dp^{1/d}+d-1$. It is also easily seen that $\underline{Size}(p;b) = {\binom{n+d-1}{d-1}}$ where $n = p^{1/d}$, so that $p^{1-1/d} = O(\underline{Size}(p;b))$.
- (3.5) Define $b: \mathbb{N}^{d} \to \mathbb{N}$ by $b(\pi) = \Pi(\pi)$. Obviously, $\beta(p; b) = p$ so that $p \leq \tau(p; b) < 2p$. (By definition $\sigma(p; b) < p$ for any b.) Moreover, by considering the cases when p is a power of 2, one verifies easily that $(\log p)^{d-1} = O(\text{Size}(p; b)).$

Note that in all of the examples just presented, the product $\tau(p; p) \cdot \underline{Size}(p; p)$ grows (in order of

magnitude) at least as fast as $p \cdot (\log p)^{d-1}$. It is our conjecture that this is a universal phenomenon; however, we are able to establish only an "infinitely often" version of this conjecture.

<u>Theorem 3.1.</u> Let $b: N^d \to N$ be an extendible hashing scheme. There is a constant c>0 such that, for infinitely many $p \in N$, $\tau(p; b) \cdot \underline{Size}(p; b) > c \cdot p \cdot (\log p)^{d-1}$.

4. EXTREME CASES

It is not uncommon that a seemingly innocuous restriction on one of our complexity measures forces another measure to behave as badly as possible (i.e., to grow as fast as one could ever force it to). In this section we present a number of such situations. To lend the reader a point of reference, we note minimax lower bounds on three of our complexity measures. (Recall that, by definition, $\sigma(p; b)$ cannot be forced to exceed p.)

- (4.1) <u>A bound on β </u>. In [2] it is proved that, if $b: \mathbb{N}^{d} \to \mathbb{N}$ is one-to-one, then $p \cdot (\log p)^{d-1} = 0(\beta(p;b))$. Moreover, for each $d \in \mathbb{N}$, there is a one-to-one scheme $b: \mathbb{N}^{d} \to \mathbb{N}$ for which $\beta(p;b) = 0(p \cdot (\log p)^{d-1})$.

We begin our survey of harmful restrictions by exhibiting two restrictions, either of which guarantees poor storage utilization.

- <u>Theorem 4.1</u>. Let b be a d-dimensional extendible hashing scheme. If either $\alpha(p;b) = 0(1)$ or $\sigma(p;b) = 0(1)$, then $p \cdot (\log p)^{d-1} = 0(\beta(p;b))$.
- Hint: The proof hinges on the fact that the array
 - scheme $(N_p)^d$ contains the union $S_d(p)$ of all p- or fewer-position arrays. See Figure 2 and Lemma 2.1.



Figure 2

Every q-position array is a subset of the $q \times q \times \cdots \times q$ array

It is worth noting that Theorem 4.1 cannot be strengthened by replacing α by $\overline{\alpha}$: in Section 5, we show that a constant $\overline{\alpha}$ can coexist with a β that is linear in p.

Expected access time must be poor for schemes that do not use increasing numbers of buckets as array size grows. Such schemes would typically use linked allocation or hashing by division (i.e., modular arithmetic as in $b(\pi) = \overline{\Sigma}(\pi)$ (mod m)).

- Theorem 4.2. Let b be an extendible hashing scheme. If $\beta(p;b) = 0(1)$, then $\log p = 0(\overline{\alpha}(p;b))$.
- Hint: Such schemes behave approximately as linkedallocation schemes (cf. (4.2)).

Finally, we note two restrictions on efficiency of storage management which, like bounded bucket demands, guarantee poor access time. The first restriction is that b assign buckets to arrays in a gap-free manner; that is, for every array A, there is an integer k for which $b(A) = \{1, \dots, k\}$. (In other, words, bucket *l-1* is used whenever bucket & is.) The second restriction, which includes special cases of bounded bucket demands, is that total storage demands grow slowly, for instance, $\tau(p;b) = p + O(p^{1/2-\epsilon})$. The fact that these restrictions give rise to worst-possible access time (in the sense discussed earlier) is somewhat more surprising than the fact that the bounded bucket restriction does. The desired results emerge as corollaries of the following theorem.

(4.3) Let b be a d-dimensional extendible hashing scheme, and let $f: \mathbb{N} \to \mathbb{N} \cup \{0\}$ be a nondecreasing function. We say that f <u>bucket-bounds</u> b if, for all $p \in \mathbb{N}$ and all array schemes $A \subset \mathbb{N}^d$ having p positions,

$$\underline{\max} \ b(A) - \#b(A) \leq f(p).$$

Theorem 4.3. Let b be a d-dimensional extendible hashing scheme. If the function f bucketbounds b, then

$$(1-1/d) \cdot \log p - \log(f(p)+1) = O(\overline{\alpha}(p;b)).$$

Hint: If f bucket-bounds b, then

 $\underline{b}(\pi) \leq (\Sigma(\pi) - d + 1) (\underline{f}(\Pi(\pi)) + 1)$

[by induction on $\Sigma(\pi)$].

- <u>Corollary 4.1</u>. If the extendible hashing scheme b_{1} is gap-free, then log $p = O(\overline{\alpha}(p;b_{1}))$.
- <u>Hint</u>: b is gap-free iff it is bucket-bounded by $f(p) \equiv 0$.
- Corollary 4.2. Let b be a d-dimensional extendible hashing scheme. If $\tau(p;b) \le p + f(p)$ where $f:N \to N\cup\{0\}$ is nondecreasing, then $(1-1/d) \cdot \log p - \log(f(p)+1) = O(\overline{\alpha}(p;b)).$
- <u>Hint</u>: The hypothesized bound on τ's rate of growth implies that f bucket-bounds b.

The interest of Corollary 4.2 for the material in this section is that it implies that a very slowly-growing τ can signal very inefficient access behavior. For instance, if $\tau(p;b) = p+0(p^c)$ for some c < 1-1/d, then log $p = O(\overline{\alpha}(p;b))$.

5. LINEARLY-GROWING BUCKET SPACE

The results in Section 4 indicate that for us even to hope to find an extendible hashing scheme that is efficient in both speed of access and utilization of storage, we must relax our stringent demands on the bucket function. Fortunately, we needn't relax our demands so much as to admit fast-

growing bucket functions: one can find efficient hashing schemes b with $\beta(p;b) = p$. For the sake of generality that will not occasion any technical encumbrance, we study here hashing schemes with $\beta(p;b) = 0(p)$ (i.e., we allow b to use more than p buckets); we call such b's <u>linear</u> hashing schemes. (Restricting attention to such schemes is really only a mild relaxation of our earlier demands since $\beta(p;b) = p$ even for the d-dimensional gap-free scheme $b(\pi) = \Sigma(\pi) - d + 1$.) The reader who balks at our allowing $\beta(p;b)$ to exceed p should note that, for any integer k, the scheme

 $b_{k}^{(k)}(\pi) = \left[b_{k}(\pi)/k\right] \text{ has the following properties.}$ (1) $\beta(p;b_{k}^{(k)}) = \left[\beta(p;b_{k})/k\right];$ (2) $b_{k}^{(k)}$ is more þ^(k) conservati

has only slightly worse access time than b (see Theorem 5.4). Therefore, we shall not be using the "big Oh" to hide large multiples that can refute our claims of efficiency.

A. Lower Bounds on Performance

The restriction to linearly-growing bucket space imposes floors on certain other efficiency criteria. We present two such lower bounds.

First we note that worst-case access time for linear hashing schemes must grow as log log p. Two points merit note. First, this bound is dramatically lower than the log p bounds on access time established in Theorem 4.2 and Corollaries 4.1

and 4.2 (this last with $f(p) = 0(p^c)$ for c<1-1/d). Second, whereas the cited results bound expected access time, the present result deals with worstcase access time.

Theorem 5.1. For any linear extendible hashing scheme b, log log $p = O(\alpha(p;b))$.

<u>Hint</u>: The array scheme $(N_p)^d$ contains the set

 $S_d(p)$ (cf. Lemma 2.1).

Our second lower bound asserts that $\ \sigma$ must grow linearly with p whenever β does. The import of this result is that, if one wishes to replace a given linear scheme by (for which $\tau(p;b) = O(p)$ automatically) by a linear scheme b' for which $\tau(p;b') = p + o(p)$, one cannot hope to do it simply by lowering interior bucket demands so that $\sigma(p;b') = o(p)$. The desired savings in storage can come only by constructing b' so as to assure that β grows slowly. (Cf. the discussion in the latter half of Section 5C on lowering storage demands.)

Theorem 5.2. For any linear extendible hashing scheme b, p = $O(\sigma(p;b))$. Specifically, if $\beta(p;b) \le c \cdot p$ for almost all p, then $\sigma(p,b) \ge \delta \cdot p$ for almost all p, where $\delta = \exp(-(c+2)).$

An Efficient Hashing Scheme в.

(5.1) For each dimensionality
$$d \in \mathbb{N} - \{1\}$$
, define the d-dimensional extendible hashing scheme $b_d: \mathbb{N}^d \to \mathbb{N}$ by the following recursion.
(a) For $\pi \in \mathbb{N}^2$, $b_2(\pi) = \pi_2 + (\pi_1 - 1)2$
(b) For $\pi \in \mathbb{N}^d$, $d > 2$,
 $b_d(\pi) = b_2(\pi_1, b_{d-1}(\pi_2, \dots, \pi_d))$.

Informally, one computes $b_{d}(\pi)$ by taking the minimum-length binary representation of π_1, \cdots, π_d , removing the leading 1 from all but that of π_1 , and concatenating the remaining strings in the appropriate order; thus the binary representation of $b_d(\pi)$ assumes the form $1\xi_1\xi_2\cdots\xi_d$, where $1\xi_1$ is the binary representation of π_{\star} .

			-	•			J			
Γ	1	2	3	4	5	6	7	8	9	10
	2	4	5	8	9	10	11	16	17	18
	3	6	7	12	13	14	15	24	25	26
	4	8	9	16	17	18	19	32	33	34
	5	10	11	20	21	22	23	40	41	42
	6	12	13	24	25	26	27	48	49	50
	7	14	15	28	29	30	31	56	57	58
	8	16	17	32	33	34	35	64	65	66
	9	18	19	36	37	38	39	72	73	74
þ	0.	20	21	40	41	42	43	80	81	82

Figure 3

A Schematic View of $b_2((N_{10})^2)$.

<u>Theorem 5.3</u>. For each dimensionalty $d \in \mathbb{N} - \{1\}$,

the extendible hashing scheme b_d has the following properties. (a) $\beta(p; b_d) = p$.

(b) There is a constant c > 0 such that

 $cp < \sigma(p; b_d) < p.$ (c) $\alpha(p; b_d) \leq$

(d-1) log log p + 0(1). (d) $\overline{\alpha}(p; b_d) = 0(1);$

in particular $\overline{\alpha}(p; b_2) < 8.5$.

<u>Proof sketch</u>. (a) First, $b_d(N_p \times \{1\}^{d-1}) = N_p$, so $\beta(p; b_d) \ge p$. On the other hand, $b_{2}(\pi) \leq \pi_{2} + (\pi_{1}-1) \cdot \pi_{2} = \pi_{1} \cdot \pi_{2}$ for all $\pi \in \mathbb{N}^{2}$; this generalizes by induction to the assertion that

 $b_d(\pi) \leq \Pi(\pi)$ for all $\pi \in \mathbb{N}^d$.

(b) The upper bound on σ is immediate by definition. The lower bound follows directly from Theorem 5.2 in conjunction with part (a).

(c) Say that $1\omega \ (\omega \in \{0,1\}^*)$ is the binary representation of $b_d(\pi)$ for some $\pi \in S_d(p)$. By part (a), $b_d(\pi) \leq p$. There are, therefore, at most $\lceil \log(p+1) \rceil^{d-1}$ ways that the string ω can be broken (or parsed) into the form $\omega = \xi_1 \cdots \xi_d$, each $\xi_i \in \{0,1\}^*$. In other words, when ξ_d stores an array having p or fewer positions, no bucket gets more than $\left[\log(p+1)\right]^{d-1}$ positions. Part (c) is now immediate by definition of Access and of α .

(d) We consider only the case d = 2. The analysis of the general case is similar. Say that b_2 stores the array scheme A of size <m,n> having p = mn position. For $z \in \mathbb{N}$, let $\ell(z) = \lfloor \log z \rfloor + 1$ denote the length of the binary representation of z. Let t = l(m) + l(n).

View the computation of b_2 as concatenation of binary representations as described in (5.1). It is clear that if bucket z is used when b_2 stores A, then $\ell(z) \le t - 1$. Consider any bucket z with $\ell(z) = t - 1$. It is not hard to see that there is at most one way z can be "parsed" (cf. part (c)) into integers π_1 and π_2 with $\pi_1 \le m$ and $\pi_2 \le n$; therefore, the population of such buckets is at most 1. There are at most 2^{t-2} such buckets. Similarly, the 2^{t-3} buckets z with $\ell(z) = t - 2$ can each be parsed in at most two ways, the 2^{t-4} with $\ell(z) = t - 3$ in at most three ways, and so on. This yields the bound $\overline{\alpha}(p;b_2) \le (1/p) \sum_{k=1}^{t+1} k \cdot \lceil \log(k+1) \rceil \cdot 2^{t-k-1}$.

Noting that $p \ge 2^{t-2}$, this sum can be bounded above by 8.5.

<u>Remark</u>. A more careful analysis in part (d) proves that $\overline{\alpha}(p; b_2) < 3$; and computer simulation reveals that $\overline{\alpha}(2500; b_2) > 2$. Of course, the average access time in practice will depend on $\overline{\alpha}$ dilated by a constant that reflects the particular data structure used to store buckets and the particular algorithm used to access positions within buckets; see [7, Section 6.2].

Even though our definition of <u>Access</u> assumes some variant of balanced search tree as the basic data structure, it is interesting to note that the average access cost of b_2 remains bounded (by 8) if one stores buckets in linear lists, charging average cost (k+1)/2 to access a position in a

average cost (k+1)/2 to access a position in a bucket of size k. Linear lists may be more attractive for "small" arrays, since the accessing algorithm is simpler for linear lists than for trees. However, the worst-case access time of b_{2} becomes roughly log p if linear lists are used.

C. <u>Tradeoffs between Access Time</u> and Storage Utilization

Theorem 5.3 does not specify precise values for $\overline{\alpha}(p; \bigcup_d)$ or $\tau(p; \bigcup_d)$, since by means of (often easily-effected) changes to a hashing scheme, one can trade off access time and storage utilization: at the cost of increasing $\beta(p; \bigcup)$ by a factor of k, one can decrease $\overline{\alpha}(p; \bigcup)$ by roughly log k; conversely, at the cost of increasing $\overline{\alpha}(p; \bigcup)$ by log k, one can decrease $\tau(p; \bigcup) - p$ by the factor k.

Lowering access time. Any change to b that causes a k-fold decrease in the size of buckets will reduce $\overline{\alpha}$ (p;b) by log k, albeit at the cost of a k-fold increase in the number of buckets used. It is difficult to discuss such alterations in the abstract, since the effectiveness of a given bucketshrinking ploy appears to depend in an essential way on the particular b in question. We limit our discussion, therefore, to a simple example in two dimensions.

Let
$$b(\pi) = \pi_1 + \pi_2 - 1$$
; this is a gap-free hash-
ing scheme (cf. Corollary 4.1). Define $b': \mathbb{N}^2 \to \mathbb{N}$

by $b'(\pi) = k \cdot b(\pi) - k + [\pi_1 \pmod{k}]$. Now compare the behaviors of both b and b' on an array $A \subset N^2$. If b assigns m array positions to a given bucket i, then b' assigns at most [m/k]positions to each of buckets $ki-k, \cdots, ki-1$. The cumulative time required to access these positions is, therefore, at most $m \lceil \log(m+1) \rceil$ for b and $m \lceil \log([m/k]+1) \rceil$ for b'. Since $\log([m/k]+1) \le$ $\log(m/k + 2) < \log(m/k) + 1$ for almost all m, our estimates on the time to access all of bucket i lead us to conclude that the average access time for A under scheme b exceeds the average access time under b' by at least $\log k - 2$ providing only that A has at least 4k rows and columns. Since

$$\alpha$$
 is monotonic in p, we see that $\alpha(p;b')$
approaches $\overline{\alpha}(p;b)$ - log k for sufficiently large
p.

Lowering storage demands. We can discuss techniques for improving storage utilization with somewhat more authority than those for improving access time, for here we find general techniques that work for all b.

(5.2) Let b_{k} be an extendible hashing scheme. For $k \in \mathbb{N}$, $b_{k}^{(k)}$ is the extendible hashing scheme defined by $b_{k}^{(k)}(\pi) = \lceil b_{k}(\pi)/k \rceil$.

If the linear hashing scheme b has total storage demand $\tau(\mathbf{p}; b) = (1+c)\mathbf{p}$ for c > 0, then the (again linear) scheme b ^(k) has demand not exceeding $(1+c/k)\mathbf{p} + 1$. Therefore, by choosing k sufficiently large, one can reduce storage demand below $(1+\epsilon)\mathbf{p}$ for any $\epsilon > 0$ by replacing b with $b^{(k)}$. Such a replacement has an adverse effect on access efficiency, but only an additive one.

Theorem 5.4. Let b be an extendible hashing
scheme. (a)
$$\tau(p;b^{(k)})-p \le 1+(\tau(p;b)-p)/k$$
.
(b) $\overline{\alpha}(p;b^{(k)}) \le \overline{\alpha}(p;b) + \log k + 1$.

Hint: Consider how b^(k) coalesces b's bucket assignments.

The following important corollary of Theorem 5.4 is immediate.

Corollary 5.1. If
$$\overline{\alpha}(p; b) = 0(1)$$
, then the same is
true of $\overline{\alpha}(p; b^{(k)})$ for any k.

We now have linear hashing schemes whose storage requirements grow as $(1+\varepsilon)p$ and whose expected and worst-case time grow as a constant and log log p, respectively. (The last remark is immediate and left to the reader for verification.) The only material improvement one might hope for is to reduce storage demands to p+o(p) without excessive deterioration of access time. We have very little to report with regard to such improvements.

Proposition 5.1. Let b be a linear hashing scheme.

The hashing scheme b_{π}^{*} defined by $b_{\pi}^{*}(\pi) =$

 $\left[b(\pi) / \log \Pi(\pi) \right]$ has the following properties.

(a) $\tau(p; b') - p \le 1 + (\tau(p; b) - p) / \log p$.

(b) $\alpha(p; b^*) = O(\alpha(p; b) + \log \log p)$.

If one is willing to use only worst-case behavior to measure efficiency of access, Proposition 5.1 yields a technique for attaining both good storage utilization and good access characteristics. However, this technique inevitably destroys efficient expected access time.

<u>Proposition 5.2</u>. For any extendible hashing scheme $b_{i}, \overline{\alpha}(p;b) \ge \log(p/\beta(p;b))$.

This proposition indicates that, for any linear $b, \log \log p = 0(\overline{\alpha}(p; p^*)) - \text{since, for some } c > 0$ $\beta(p; p^*) \leq \left\lceil cp/\log p \right\rceil.$

We are thus left with the following open question.

<u>Problem</u>. Does there exist an extendible hashing scheme b such that $\tau(p;b) = p+o(p)$ and $\overline{\alpha}(p;b) = 0(1)$?

Our results do not resolve this issue, but they do narrow down the search materially. Specifically we know that the desired b (if it exists) must satisfy the following efficiency bounds.

(5.5) (a) For some k,k' > 0, for almost all p,

 $kp \leq \beta(p;b) \leq k'p.$

(b) For some
$$l > 0$$
, for almost all p

 $lp \leq \sigma(p;b) < p.$

The upper bound on β follows from our desire to keep storage demands low; the lower bound follows from Proposition 5.2 and our desire to keep expected access time down. The upper bound on σ follows by definition; the lower bound on σ follows from Theorem 5.2 in the presence of the linearity of b, i.e., (5.5(a)).

Another remaining challenge, of a materially different nature, is to "fine tune" the results in this section in an actual computer environment, that is, select the necessary data representations, investigate the cost of bookkeeping details that are not visible in any abstract investigation, and study the real effect of the tradeoffs mentioned. We suspect, on the basis of Theorems 5.3 and 5.4 and the remarks

on lowering $\overline{\alpha}$, that extendible hashing schemes that attain a practical level of efficiency exist.

ACKNOWLEDGMENT

We are grateful to Shmuel Winograd for constructive suggestions in the early stages of this research.

REFERENCES

- A. L. Rosenberg, Allocating storage for extendible arrays. JACM, <u>21</u> (1974) 652-670.
- A. L. Rosenberg, Managing storage for extendible arrays. <u>SIAM J. Comput.</u>, to appear.
- 3. D. E. Knuth, <u>The Art of Computer Programming I:</u> <u>Fundamental Algorithms</u>, Addison-Wesley, <u>Reading</u>, Mass., 1968.
- 4. E.v.d.S. de Villiers and L. B. Wilson, Hash coding methods for sparse matrices. Tech. Rpt. 45, Univ. of Newcastle-upon-Tyne (Computing Lab.), May, 1973. See also, Hashing the subscripts of a sparse matrix. <u>BIT</u>, <u>14</u> (1974) 347-358.
- L. J. Stockmeyer, Extendible array realizations with additive traversal. IBM Report RC-4578, 1973.
- A. L. Rosenberg, Computed access in ragged arrays. in <u>Information Processing 74</u> (J. Rosenfeld, ed.) North-Holland, Amsterdam, 1974, pp. 642-646.
- 7. D. E. Knuth, <u>The Art of Computer Programming III:</u> Sorting and <u>Searching</u>, Addison-Wesley, Reading, Mass., 1973.
- O. Amble and D. E. Knuth, Ordered hash tables. <u>Comput. J.</u>, <u>17</u> (1974) 135-142.