

Compression of Correlated Bit-Vectors *

A. Bookstein and S.T. Klein

Center for Information and Language Studies
University of Chicago, 1100 East 57-th Street
Chicago, IL 60637

July 1990

Abstract:

Bitmaps are data structures occurring often in information retrieval. They are useful; they are also large and expensive to store. For this reason, considerable effort has been devoted to finding techniques for compressing them. These techniques are most effective for sparse bitmaps. We propose a preprocessing stage, in which bitmaps are first clustered and the clusters used to transform their member bitmaps into sparser ones, that can be more effectively compressed. The clustering method efficiently generates a graph structure on the bitmaps. In some situations, it is desired to impose restrictions on the graph; finding the optimal graph satisfying these restrictions is shown to be NP-complete. The results of applying our algorithm to the Bible is presented: for some sets of bitmaps, our method almost doubled the compression savings.

1. Introduction

Textual Information Retrieval Systems (IRS) are voracious consumers of computer storage resources. Most conspicuous, of course, is the text itself, which constitutes the content of the database. But, to efficiently use the database, auxiliary structures must be created that themselves require a substantial amount of space. Thus, mechanisms for compressing a wide range of data structures must be sought for the efficient operation of such systems [10]. To date, most attention has been given to, and progress made in, the area of text compression ([2], [13], [16]). In this paper, we shall describe and examine the possibilities of compressing bitmaps, a data structure often proposed for improving the performance of retrieval systems ([6], [18]).

Bitmaps occur often in information retrieval. They can represent the occurrences of a word in the sentences or paragraphs making up a text; they can indicate the documents associated with an index term; they appear as bit slices of a matrix of signatures; they might represent pixels in rows of a raster graphics display. They are useful; they are also large and expensive to store. Much work has been carried out on

* This paper extends the work presented at the 13-th ACM-SIGIR Conference (Brussels, September 1990) under the title *Construction of Optimal Graphs for Bit-Vector Compression*.

the compression of bitmaps, and this has been especially successful for those that are very sparse. But not all bitmaps are sparse, and even sparse bitmaps could benefit from further compression. This paper describes a method that complements existing compression techniques and improves their performance, at least for certain categories of bitmaps.

We concentrate on sets of bitmaps such as are generally found in information retrieval (IR) systems. How such bitmaps can be used to enhance the system is discussed in [5] and [3]. Each bit-position corresponds to a specified sub-unit of the database, henceforth referred to as a *segment*; below, a segment will refer to a paragraph of text, though, in other contexts, a full document (or even a set of documents) may be the preferred unit. For each different word (or index term) W in the database, there is a map $B(W)$, such that the i -th bit of $B(W)$ is 1 if and only if W appears in (or is assigned to) segment i . Such bitmaps can be compressed very efficiently. In part this is because they tend to be very sparse. That bitmaps compress better as they become sparser is expected theoretically. For suppose a bitmap can be considered as having been produced by a random bit generator, with the probability of a one bit being p (the theory can easily be extended to encompass more complex models of bitmap generation). Then the information content of a bit is given by:

$$H = -p \log p - (1 - p) \log(1 - p),$$

and, for a bitmap of ℓ bits, the quantity ℓH forms a lower bound on the number of bits needed to represent the bitmap. As is well known, H increases monotonically as p increases from 0 to .5, and then decreases monotonically as p continues growing. Since almost all of our bitmaps have p less than .5, we expect compression to improve as p decreases, that is, as the map becomes sparser. (For $p > .5$, we could complement the bitmap before proceeding.) Thus we wish to be alert to opportunities for reducing the density of our bitmaps; this is the essence of the approach described in this paper.

Other factors also contribute to our ability to compress bitmaps effectively, as evidenced by the fact that actual IR bitmaps are more compressible than randomly generated bitmaps with the same density of 1-bits [4]. The reason for the better results is a *cluster-effect*: since the segment positions in the bitmaps are usually ordered by topic or chronologically, adjacent bits often correspond to segments treating the same or related subjects. Thus the appearance of a word in a given segment often implies that it also appears in neighboring segments. This effect is exploited by many compression methods, resulting in excellent reduction in storage requirements.

There is, however, another clustering possibility that has hitherto been overlooked, one involving sets of bitmaps (words), rather than sets of bits (segments) within a single bitmap. The occurrences of certain words, especially those taking part in well known phrases like **Security Council** or **Curriculum vitæ**, are sometimes strongly

correlated across segments in the sense that if one word appears in a certain segment, the other is also very likely to do so. Such pairs of bitmaps are likely to be quite similar. But identifying clusters of such highly associated words is not as direct as it was for bit clusters within a bitmap, because words and their associated bitmaps are generally arranged in lexicographical order, not in order of logical proximity. In this respect, IR bitmaps differ from self-clustered graphic bitmaps, in which adjacent (raster) rows are often similar. The objective of this paper is to show how to usefully identify clusters of correlated words, and then take advantage of these associations to squeeze out some additional compression.

In Section 2, we briefly review some known bitmap compression techniques and propose a new one that is simple and easy to implement; it will then be used as the compression component of a two-stage compression process described in Section 3. The first stage of the two stage process is to partition the bitmaps of our IR system into clusters of correlated bitmaps; the resulting clusters are then used to transform the original bitmaps into another set of bitmaps that are sparser and more effectively compressed in stage 2 of the process. In Section 4 we report on experiments testing the new method; the database we chose to study was the Hebrew Bible. We then consider in Section 5 variants that address some problems raised by the example, and would on occasion be useful; a couple of related problems are then shown to be NP-complete.

2. Bit-vector compression techniques

2.1 Overview of some known methods

Suppose we are given a bitmap v of length ℓ bits, of which s are ones and $\ell - s$ are zeros. In our applications, the maps are usually sparse, i.e., $s \ll \ell$. The simplest way to store v compactly for very small n is to **enumerate the positions** of the 1-bits. As one needs $d = \lceil \log_2 \ell \rceil$ bits to identify any position, this method would need sd bits for each map, which may well be much smaller than the ℓ bits required for the uncompressed original map. Alternatively, one could record the distances between successive 1-bits, that is, give the position of a 1-bit relative to the preceding 1-bit position rather than relative to the beginning of the vector. This is known as **run-length coding** (Schuegraf [14]). In its simplest form, the length of every run is encoded by a fixed length codeword; since this codeword must be large enough to accommodate the theoretical maximum run length, this is equivalent to the previous method.

Since, in simple run length coding, the space allocated for each run must be adequate for the largest possible run, such codes can be inefficient if many of the runs are of small or moderate length. The following variant, due to Teuhola [17], improves on simple run length coding by having a **variable length** representation of

a **run length**. A run of r zeros is first broken up into successive blocks of zeros of exponentially increasing size; the first block is of size 2^k (for k a parameter selected to optimize this procedure), the second of size 2^{k+1} , etc., until a block is produced that extends beyond the run, i.e., is partially filled. The length of the run, r , can then be represented as follows: (a) each block in the sequence that is completely filled with zeros is represented in turn by a one, and (b) a zero is appended to the string of ones as delimiter. If t ones are present, then we know $2^k + 2^{k+1} + \dots + 2^{k+t-1} \leq r < 2^k + \dots + 2^{k+t-1} + 2^{k+t}$, that is: the first t blocks are filled with zeros, but the last potential block of 2^{k+t} zeros is either empty or partially filled. So, finally, (c) we can explicitly represent the number of zeros in the last block as a binary integer with $k+t$ bits. Thus a run of length r is encoded by $O(\log r)$ bits instead of $O(\log(\text{max length}))$.

Jakobsson [9] suggests the use of **Huffman coding** for bitmaps. The bit-vector is partitioned into blocks of fixed size k , and statistics are collected on the frequency of occurrence of the 2^k bit patterns. Based on these statistics, the set of blocks is Huffman encoded, and the bitmap itself is encoded as a sequence of such codewords. For sparse vectors, the k -bit block consisting of zeros only, and blocks with only a single 1-bit, have much higher probabilities than the other blocks, so the average codeword length of the Huffman code will be smaller than k .

Fraenkel & Klein [7] combine **Huffman coding with run-length coding**. Once again, a parameter k is chosen as a block size. However, since for very sparse vectors the probability of a block of k zeros is high, runs of blocks of k zeros receive special treatment. We first represent the succession of k -bit blocks comprising a bitmap as a sequence of two categories of symbols: beginning with the first block, if a k -bit block includes 1-bits, then we represent it by its own special symbol, as in the previous method. If it is a zero-block, then instead of representing the block itself, we represent the entire run of zero blocks which it starts by a string of integers as follows: suppose the run consists of r zero-blocks, with r represented in binary form as $r = \sum_{i \geq 0} a_i 2^i$, for a_i zero or one. Then the run is represented in the symbol sequence by the string of integers n_0, n_1, \dots , where each n_i is a power of 2 in the representation of r for which $a_{n_i} = 1$; this in effect encodes the run lengths. Next, the frequency of occurrence throughout the bitmap file of each special and integer symbol is recorded, permitting a Huffman tree to be constructed for the $2^k - 1$ special symbols together with the integer symbols. Finally, the bitmap is Huffman encoded using this tree.

A **hierarchical method** for compressing a sparse bitmap was proposed by Wedekind & Härder [19]. The original bit-vector v_0 of length ℓ_0 bits is partitioned into r_0 equal blocks of k_0 bits each ($r_0 \cdot k_0 = \ell_0$), and the blocks consisting only of zeros are dropped. The resulting sequence of non-zero blocks does not by itself allow the reconstruction of v_0 ; we can, however, append a list of the indices indicating where these non-zero blocks occur in the original vector. This list of up to r_0 indices is itself

kept as a bit-vector v_1 of $\ell_1 = r_0$ bits; there is a 1 in position i of v_1 if and only if the i -th block of v_0 is not all zeros. Now v_1 can be further compressed by the same method. In other words, a sequence of bit-vectors v_j is constructed, each bit in v_j being the result of ORing the bits in the corresponding block in v_{j-1} . The procedure is repeated recursively until a level is reached where the vector length reduces to a few bytes. The compressed form of v_0 is then obtained by concatenating, in order of decreasing i , all the nonzero blocks of the various v_i . The same method appears in Vallarino [18], who used it for two-dimensional bitmaps, but with only one level of compression.

The hierarchical method is refined by Choueka & al. [4], by adding a **pruning algorithm** that removes from the hierarchy-tree, branches pointing to very few segments. The algorithm partitions the set of 1-bits in v_0 into two subsets: the class of 1-bits which are efficiently handled by the hierarchical method; and the complementary class, consisting of more or less isolated 1-bits whose inclusion in the hierarchical tree structure would have been more expensive than their enumeration in an appended list. Either of these two classes may be empty. If the list is long enough, it is further compressed by a variant of prefix omission, to be described in more detail in the following sub-section.

It should be noted that since, for each map, the number of runs of zeros is equal to the number of 1-bits plus 1, the size of the compressed file obtained by the first few methods is clearly linearly related to the number of 1-bits in the original file. For the hierarchical and Huffman coding methods this relation is less evident, but has been empirically established. This observation is consistent with the theoretical argument presented in the introduction, and reinforces our intention to design a preprocessing stage that reduces the number of bits in a bitmap.

2.2 A simple new method

The following technique is a simple generalization of the prefix omission method suggested in [4] for the secondary compression of the list of 1-bits which were pruned from the tree. It can also be viewed as a variant of the hierarchical method, using only a single level of compression.

Choose an integer parameter k and partition the original vector v_0 of length ℓ_0 into blocks of 2^k bits. We shall assume that the number of 1-bits in the bitmap is s . As in the hierarchical method, construct a new vector v_1 of length $\lceil \ell_0/2^k \rceil$, in which bit i is zero if and only if block i of v_0 contains only zeros. However, now, instead of storing the non-zero blocks of v_0 themselves, we substitute for each block the string of indices of the 1-bits within that block. *A priori* k bits are sufficient for storing such a relative index; however we need an additional bit per index to serve as a flag, which identifies the boundary of each block. Therefore, in addition to the fixed overhead of

storing the vector v_1 , $k + 1$ bits are needed for representing each of the 1-bits of v_0 . When a block has a small number of 1-bits, a significant saving in space could result.

We would now like to find k that optimizes the size of the block to be chosen, that is, the integer k^* that minimizes $f(k) = \lceil \ell_0 2^{-k} \rceil + (k + 1)s$, the size in bits of the compressed bitmap. Because of the appearance of the ceiling function in $f(k)$, finding the minimum value directly is difficult. Instead we shall search for an integer k_1^* that minimizes the related continuous function $f_1(k) = \ell_0 2^{-k} + (k + 1)s$. Since f_1 is a convex function, k_1^* is an integer which satisfies

$$f_1(k_1^*) < f_1(k_1^* + 1) \quad \text{and} \quad f_1(k_1^*) \leq f_1(k_1^* - 1).$$

Examining the left hand inequality, we find

$$\ell_0 2^{-k_1^*} + (k_1^* + 1)s < \ell_0 2^{-k_1^*+1} + (k_1^* + 2)s,$$

or, $2^{-k_1^*+1} < s/\ell_0$. Thus $k_1^* > \log_2(\ell_0/s) - 1$. Similarly, the right hand inequality is equivalent to $k_1^* \leq \log_2(\ell_0/s)$. Combining the two, we find that k_1^* must satisfy

$$\log_2 \frac{\ell_0}{s} - 1 < k_1^* \leq \log_2 \frac{\ell_0}{s},$$

so $k_1^* = \lfloor \log_2(\ell_0/s) \rfloor$.

If we have a file of m bitmaps, we want to use the same method for encoding each of them, so the optimal k will be determined by the average \bar{s} of 1-bits per map. The total number of bits in the compressed file is thus

$$m \left[\ell_0 2^{-\lfloor \log_2(\ell_0/\bar{s}) \rfloor} \right] + S(\lfloor \log_2(\ell_0/\bar{s}) \rfloor + 1), \quad (1)$$


for $S = m\bar{s}$, the total number of 1-bits in the bitmap set.

A priori, k_1^* need not equal k^* ; however, it is easy to see that the cost of using k_1^* is identical to the cost of using k^* . To show this, first note that by the definition of the ceiling function, $f(k_1^*) < 1 + f_1(k_1^*)$. Since k_1^* minimizes f_1 , $f_1(k_1^*) \leq f_1(k^*)$. But f_1 cannot exceed f for any k , and in particular $f_1(k^*) \leq f(k^*)$. Combining these results with the fact that k^* minimizes f , we get

$$f(k^*) \leq f(k_1^*) < 1 + f(k^*).$$

We conclude that using k_1^* for the true optimum, k^* , results in an excess of less than one bit in storage for each bitmap. But $f(k)$ takes only integer values at integer k , so if the difference $f(k_1^*) - f(k^*)$ is smaller than 1, then it must actually equal zero: either $k_1^* = k^*$, or, at least, the storage implications are the same for both values ($f(k_1^*) = f(k^*)$).

For example, suppose $\ell_0 = 180$ and the indices of the 1-bits in v_0 are 36, 50, 53, 105 and 126. Thus $s = 5$, so we get that the optimal k is $\lfloor \log_2(180/5) \rfloor = 5$. There are $\lceil 180/2^5 \rceil = 6$ bits in v_1 , each (except the last) corresponding to a block of 32 bits in v_0 . There are three 1-bits in the second block, with relative indices 4, 18 and 21, and there are two 1-bits in the fourth block, with relative indices 9 and 30; the four other blocks are empty. Thus the following information would be kept:



The number of bits necessary to store this map is thus $6 + 5 \times (5 + 1) = 36$. With $k = 4$ we would need $12 + 5 \times (4 + 1) = 37$ bits and with $k = 6$ we would need $3 + 5 \times (6 + 1) = 38$ bits. Note that if we list the relative indices of each sub-range in increasing order, the flag identifying the last index of each range is not always needed. In our example, for instance, the list of stored relative indices is 4, 18, 21, 9, 30, so clearly the sublist corresponding to the second 1-bit in v_1 consists of the last two elements. If, however, there were no 1-bit in position 105 of v_0 , the list of stored relative indices would have been 4, 18, 21, 30, and the partition of this list into two increasing sub-lists is not uniquely determined.

3. Bitmap Clustering

3.1 Motivation

We have remarked several times above that sparser bitmaps are more effectively compressed. We will now describe a method for reducing the number of 1-bits by making use of a natural clustering of bitmaps. To do this, we take advantage of the fact that many bitmaps are associated in the sense that the presence of a 1-bit in one map increases the likelihood of a 1-bit occurring in the same position in the other. If two bitmaps X_1 and X_2 are strongly associated in this sense, then the bitmap $X_3 = X_1 \text{ XOR } X_2$ will very possibly have fewer 1-bits than, say, X_2 . If we store X_1 and X_3 , we can reconstruct X_2 . The advantage of doing this is that we may be able to compress X_1 and X_3 more effectively than the original vectors. Since our intention when XORing two vectors is to reduce the number of 1-bits, it is useful to take as a measure of association between two vectors, the number of 1-bits in the XORed vector. But this quantity is the familiar *Hamming distance* between the two vectors. If the maps X_i and X_j are “close” in the Hamming distance sense, we would want to keep \bar{X}_j and the pair (\bar{x}_i, j) instead of \bar{X}_j and \bar{X}_i ; here $x_i = X_i \text{ XOR } X_j$, and a bar indicates that the maps have been compressed, say by the method presented in Section 2.2. Given the retained information, the original bitmap can then be recovered

by first decompressing \bar{x}_i and \bar{X}_j , which yields x_i and X_j , and finally XORing again, since $X_i = x_i \text{ XOR } X_j$.

As described above, the unchanged map X_j is compressed directly. However, X_j may itself be quite close to a third map, X_k , and therefore profitably XORed with that third map, producing the pair (\bar{x}_j, k) . Continuing in this manner we impose a structure on the bitmaps that can be represented as a directed graph, $G = (V, E)$, where the vertices $V = \{X_1, \dots, X_m\}$ correspond to the bitmaps and (X_i, X_j) , the directed edge from X_i to X_j , belongs to E if and only if X_i is compressed as (\bar{x}_i, j) .

To be workable, the following restrictions must be imposed on G . **(1)** Any map can be compressed by XORing with at most one other map, so *the outdegree of every vertex is at most 1*. **(2)** In a general graph satisfying condition (1), it might be possible to form a chain of bitmaps $X_1, X_2, \dots, X_k, X_1$, denoting that X_i is stored as $(\bar{x}_i, i + 1)$ for $i = 1, \dots, k - 1$, and X_k is stored as $(\bar{x}_k, 1)$. However, this situation must be prohibited, if we want to be able to recover the original bitmaps: starting with an arbitrary node, the chain must terminate with an untransformed bitmap, that is, with a node with outdegree zero. In other words, *a legitimate graph must be cycle free*.

These conditions impose a strong structure on a legitimate graph. Let $R = \{r_1, \dots, r_n\}$ be the set of vertices with outdegree zero, and define $\mathcal{T}(r_i)$ as the set of vertices from which there is a directed path to r_i ; $\mathcal{T}(r_i)$ also includes r_i (connected to itself by the empty path). Since there are no cycles in G , a directed path starting at any vertex $X \in V$ must eventually terminate, reaching one of the vertices $r_i \in R$. Thus every $X \in V$ is in one of the $\mathcal{T}(r_i)$. If $X \in \mathcal{T}(r_i) \cap \mathcal{T}(r_j)$ for $i \neq j$, then some vertex in the chain starting at X must have outdegree ≥ 2 . Since this is impossible, the components $\mathcal{T}(r_i)$ are disjoint and $\{\mathcal{T}(r_i)\}$ is a partition of V into connected *clusters* of bitmaps. Further, there is no linkage between any pair $\mathcal{T}(r_i)$ and $\mathcal{T}(r_j)$: for suppose $X_1 \in \mathcal{T}(r_i)$ and an edge (X_1, X_2) exists with $X_2 \in \mathcal{T}(r_j)$. But then, since a path exists connecting X_2 to r_j , a path exists (through X_2) connecting X_1 to r_j . Such a node X_1 is a member of both $\mathcal{T}(r_i)$ and $\mathcal{T}(r_j)$, which is impossible. The $\mathcal{T}(r_i)$ are thus isolated connected components in G ; because of conditions (1) and (2), each $\mathcal{T}(r_i)$ is an *oriented tree*, as defined by Knuth [11, Section 2.3.4.2].

Any forest of bitmaps can serve as the basis of our precompression operations. To maximize compressibility, however, we want to choose that forest among all possible forests that minimizes the total number of ones in the resulting bitmaps. (There could conceivably exist some maps which, because of their special internal structure, yield better compression than others which are sparser. But until a quantitative relationship can be derived between detailed bitmap characteristics and compression, sparseness is the best measure we have for bitmap compressibility.) We define the quantity to be minimized, that is, the total number of 1-bits in the roots plus the total number of 1-bits in the XORed bitmaps, as the *cost* C of the forest. Note that adopting this criterion prevents our XORing two vectors when the result would increase the number of 1-bits — for example, in the extreme case, the set of original

bitmaps, with no XORed maps, is a forest and thus a legitimate graph.

An exhaustive search generating all the possible graphs satisfying our constraints and checking for each the cost for the forest, must be ruled out on the grounds of computational expense, even if we have only a moderately large number m of bitmaps: the number of ways to partition the m maps into n clusters is a Stirling number of the second kind $\left\{ \begin{smallmatrix} m \\ n \end{smallmatrix} \right\}$ (see [11, Exercise 1.2.7–64]). But $\left\{ \begin{smallmatrix} m \\ n \end{smallmatrix} \right\}$ is asymptotically $e^n n^{m-n-1/2}$ [1, Section 24.1.4], so summing over the potential number of clusters, we find that the number of partitions is $\sum_{n=1}^m \left\{ \begin{smallmatrix} m \\ n \end{smallmatrix} \right\} > \left\{ \begin{smallmatrix} m \\ m/2 \end{smallmatrix} \right\} = \Omega(m^{m/2})$. Moreover, to get the number of possible graphs to be checked, we must consider, for each partition, the number of ways to organize each cluster internally as an oriented tree.

Fortunately, such a search is unnecessary, as the problem is equivalent to another for which there are well known polynomial algorithms. To see this, we first recall that, except for the roots, the number of 1's in a XORed bitmap is just the Hamming distance between it and its successor in the directed graph. Thus, if we assign this distance as a weight to each edge, the cost of a forest is simply the sum of the weights of all edges in the graph plus the sum of the number of ones in each root. But we can further simplify the statement of the problem by noting that the number of ones in a map is its Hamming distance to the zero bitmap (the bitmap, all of whose values are zero), denoted by X_0 . Thus, given any forest, if we introduce the zero bitmap and include the weighted edge between each root and X_0 (thereby transforming the forest into a tree), then the cost of the original forest is just equal to the sum of edge-weights over all the edges of the resulting tree in the enhanced graph. The latter sum will be called the *cost* of the tree. Since only the weights are significant when computing the cost, we can consider the tree as being non-directed. Such a simplification is well defined since the weight on an edge does not depend on its orientation (the Hamming distance is a symmetric measure). Thus given any directed forest over the set of nodes V , we can define a non-directed tree over $V^* = V \cup \{X_0\}$ having the same cost.

The converse is also true. First note that given the set of vertices V^* , any (undirected) spanning tree on V^* defines a directed tree on V^* : the root of the directed tree is X_0 ; the directed edge (X_i, X_j) is in the directed tree if a path (X_i, X_j, \dots, X_0) exists in the undirected tree. Next, by removing X_0 and the edges incident on it from the directed tree, we obtain a directed forest G , on V . Furthermore, if the edge weights are as defined above, the cost of the forest is equal to the cost of the tree that induced it. Because of this equivalence, an optimal forest is associated with an optimal (lowest cost) tree. Thus our problem is equivalent to the following one: given a complete non-directed graph whose vertex set is the union of our bitmaps with X_0 , and for which the weight on edge (i, j) is the Hamming distance between vertices i and j , find the tree for which the total edge weight is minimum. The directed forest induced by this tree is the solution to our problem.

More formally: we are looking for a graph G , which is a forest of oriented trees

spanning the vertex set V , optimizing our problem. To find the graph G , we consider the weighted undirected graph $G^* = (V^*, E^*)$, where the set of vertices V^* is obtained by adjoining a new vertex, the zero vector X_0 , to the set V of bitmaps; $E^* = V^* \times V^* - \{(X_i, X_i) : X_i \in V^*\}$ (ignoring order), that is, G^* is a complete graph from which self-loops are removed; and the weight $w(i, j)$ associated with the edge $(X_i, X_j) \in E^*$ is the Hamming distance between X_i and X_j . We then define as a legitimate sub-graph of G^* a non-directed tree T connecting all the vertices in V^* . Our task is to find the legitimate sub-graph for which the sum of all the weights of the edges in T is minimized — in fact, a minimum spanning tree (MST) of G^* . The MST in G^* now induces the optimal directed forest, G , on the original set of bitmaps, as described above. The vertices that were adjacent to vertex X_0 in T are the roots of the oriented trees in G . G is the optimal forest we were seeking.

Many algorithms for finding a MST for a non-directed graph appear in the literature, ranging from Kruskal’s simple greedy algorithm [12], which has in our case complexity $O(m^2 \log m)$, to Yao’s more involved technique [20], which would need $O(m^2 \log \log m)$ operations for our application.

3.2 Algorithm statement

Summarizing, we suggest the following procedure as the first stage for compressing a set of m bitmaps X_1, \dots, X_m . This method in principle improves any given compression algorithm \mathcal{C} for individual bitmaps, provided our assumption of strong correlation between some of the maps holds. As output, we get a table B of compressed bitmaps, the compressed form of X_i being stored in $B(i)$, $1 \leq i \leq m$. In addition, the algorithm produces a small table F of size m , defined by $F(i) = j$ if the map X_i is compressed as (\bar{x}_i, j) (i.e., if X_j is the father of X_i in the oriented rooted tree T), or by $F(i) = 0$, if X_i is the root of one of the trees.

1. Choose a compression method \mathcal{C} for an individual map: given a bitmap X , $\mathcal{C}(X)$ is the result of \mathcal{C} applied to X .
2. Extend the set of bitmaps by adjoining X_0 , the zero-vector.
3. (a) Using the Hamming distances as weights on the complete graph without self-loops having $\{X_0, X_1, \dots, X_m\}$ as set of vertices, compute a minimum spanning tree T .
 (b) Consider T as an oriented tree rooted at X_0 .
 (c) The subtrees of X_0 in T partition the original set of bitmaps.
4. (a) If X_i is a vertex adjacent to X_0 in T , then it is the root of one of the oriented trees: these bitmaps (one per tree) are compressed directly using \mathcal{C} .

$$B(i) \leftarrow \mathcal{C}(X_i)$$

$$F(i) \leftarrow 0$$

- (b) A bitmap X_i , which is not the root of a tree, has a directed edge to

another bitmap X_j in the same tree; X_i is compressed by first computing $x_i = X_i \text{ XOR } X_j$ and then compressing x_i using \mathcal{C} .

$$B(i) \leftarrow \mathcal{C}(x_i)$$

$$F(i) \leftarrow j$$

In the case of a set of bitmaps of an IR system, the problems of compression and decompression are not exactly symmetric. Compression is performed only once, during the construction of the system, and is applied to the entire set. Decompression, on the other hand, is practically never needed simultaneously for the entire set, but only for those maps associated with the keywords of a submitted query. We thus present the procedure `decompress(i)` which returns the original map X_i . It uses the function \mathcal{C}^{-1} as the inverse of the compression function \mathcal{C} — that is \mathcal{C}^{-1} decompresses bitmaps which have been compressed by \mathcal{C} .

```

decompress( $i$ )
  if  $i = 0$    return(  $\mathcal{C}^{-1}(B(i))$  )
  else       return(  $\mathcal{C}^{-1}(B(i)) \text{ XOR } \text{decompress}(F(i))$  )

```

We see that the savings in storage space gained by our clustering procedure come at the expense of increased processing time. In order to recover the bitmap X_i , we need decompress all the maps forming the path from X_i to the root of the cluster X_i belongs to.

4. Example

The database we chose for testing our algorithm is the Hebrew Bible, consisting of 305514 words which are partitioned into 929 chapters. The number of different words is 39647. Following the suggestion in [5] that bitmaps should be constructed only for words which appear more often than some fixed frequency threshold, we restricted ourselves to the 1478 words which appeared in at least 20 chapters. As a text segment, we defined a set of four consecutive chapters. The resulting bitmaps were $\lceil 929/4 \rceil = 233$ bits long. The total number of 1-bits in the 1478 maps was 65734, or $\bar{s} = 44.47$ 1-bits per map.

We first used the compression technique of Section 2.2 by itself. The optimal parameter k was $\lfloor \log_2(233/44.47) \rfloor = 2$. From equation (1) we thus get that the total number of bits needed to store the set of bitmaps in compressed form is 284404. For the uncompressed file we would need $1478 \times 233 = 344374$ bits, so that the simple method yields 17.4% compression. This k is indeed optimal for this method, since with $k = 1$ we get 11.6% compression, and with $k = 3$ we get 10.8%.

We then applied Kruskal's MST algorithm, which partitioned the set of bitmaps into 716 clusters. Of these, 530 were singletons, i.e., maps which couldn't effectively be XORed with some other map and which were therefore compressed without transformation. The other 948 bitmaps were partitioned into 186 clusters, each containing

at least two elements. Since in each cluster, the root is compressed directly, the number of bitmaps which were XORed before compression was $948 - 186 = 762$. For these, the total number of 1-bits decreased from 48590 to 33538, that is, by 31%.

Considering the entire file of bitmaps, the overall number of 1-bits decreased from 65734 to 50658, or to $\bar{s} = 34.27$ 1-bits per bitmap. The optimal parameter k was thus $\lfloor \log_2(233/34.27) \rfloor = 2$, as before. Substituting the values for m , S and \bar{s} in equation (1), we find that the total number of bits needed to store the set of bitmaps if we use the clustering method of Section 3 is 239176. Relative to the noncompressed file this is a 30.5% reduction, and relative to using only the method of Section 2.2 without clustering, this is a 15.9% improvement.

It is interesting to compare this to the information theoretic estimate of compressibility mentioned in the introduction. The probability of a 1-bit in the original file is 0.19, yielding an entropy per bit of $H = 0.703$. This means that if the 1-bits appear with the given frequency but independently from each other, the best possible compression would be 29.7%. Indeed, we got only 17.4% when the bitmaps were compressed individually. Introducing the clustering, we exploit the dependencies between different bitmaps, yielding compression savings of 30.5%, which is beyond those possible for independently generated maps.

While most of the generated clusters were small (two to four elements), some formed deep trees with tens of bitmaps, and the largest consisted of a tree of depth 15 with 112 vertices. A closer look at some of the larger clusters revealed interesting associations. Figure 1 shows a typical example. For each node in the tree, the Hebrew word is first given in English transliteration, using $\{\mathbf{ABGDHWZXtYKLMNSaPCQR\$T}\}$ respectively for $\{\textit{aleph, beth, \dots, tav}\}$, as well as the translation of the word into English.

In this cluster, 18 out of 28 words are numerals; these are clearly connected, as the Bible tends to give exact dimensions (note the words `length` and `cubit`) in certain detailed descriptions. See, for instance, Exodus 27:9–19, where a description of the `court`, the root of this cluster, is given. The depth of this tree is 5, which is therefore the maximal depth of the recursion for the decompression algorithm. Note also that the root of this cluster has a high in-degree. This was not always the case, as can be seen in the following example.

To present the second example, we use a more compact representation, based on pre-order traversal of a tree. A tree can be represented recursively by its root, followed by the list of its subtrees enclosed in parenthesis. To improve readability, the level of the root of a subtree in the full tree appears as subscripts to the parentheses. We now give only the English translation of the word at each node (many Hebrew words must be translated into several English words).

and they camped (₁ night, and they saw, and they went out (₂ the men (₃ and they came (₄ from before, and he sent (₅ and they said, and he sat (₆ and he, his people, and he went out)₆, and

HXCR
the court

Wa\$RYM
and twenty

TMYM
perfect

AMH
cubit

aMDYM
standing

HaLH
burnt offering

W\$TY
and two

ARK
length

WXM\$YM
and fifty

WXM\$H
and five

W\$L\$
and three

\$Ba
seven

HAXD
the one

W\$L\$YM
and thirty

\$MNH
eight

AXT
one

WXM\$
and five

\$TY
two

WRB
and most

\$TYM
two

XM\$H
five

ARBaH
four

a\$RH
ten

\$L\$H
three

\$S
six

HaYR
the city

A\$R
that

a\$W
they made

Figure 1: Sample cluster produced by the MST algorithm

he called (6 and he came (7 and now (8 please)8)7, and he gave
 (7 in the hand of)7, and he took (7 bread, and he did (8 two)8
)7)6)5, and they sat)4, and they went)3)2)1

The 25 elements in this cluster form a tree of depth 8, but no node has higher in-degree than 3. Note that most of the words are verbs related to motion, all in the past tense, and in third person singular or plural.

We also tried to apply the algorithm to sparser bitmaps, by defining a segment to be one, instead of four, chapters. The 1478 bitmaps then had a total of 95472 1-bits, so they were compressed with $k = 3$ and gave 59.9% compression. The clustering algorithm however produced only 300 bitmaps that were XORed; for these, the reduction in the number of 1-bits was about 21%, but nonetheless, the total number of 1-bits remained quite large at 85195 bits. The optimal k now shifted to be 4, and compression was improved by 7.5%; measured relative to the full file, a 62.9% reduction was achieved. In this case, the theoretical optimum for independently generated bitmaps with this 1-bit density is 63.8%.

In order to check the influence of the language of the database on the algorithm, we repeated the experiments with the *King James Bible*, again with a segment equal to one chapter. The improvement of the clustering method in this case was only 5.7%, again because only a small number of bitmaps (377 of 1454) were XORed. There were nevertheless some interesting clusters. For example: Asher (1 Ephraim, Joseph, Manasseh, Simeon (2 Levi, Reuben (3 Gad)3)2, Zebulun (2 Naphtali)2, Issachar (2 Benjamin, Dan)2)1. This cluster contains the names of all the tribes, except Judah. The latter appears in another cluster, together with words like Jerusalem, reign, reigned, kings, etc. Clearly, Judah differs from the other tribes, as his name often refers to the kingdom or land of Judah.

Summarizing our experiments, we see that the clustering algorithm works better when the bitmaps are not so sparse: very sparse vectors tend to have very few overlapping 1-bits, so that there is often no gain to be achieved by XORing. However, for the very sparse vectors, many of the known techniques already yield excellent results. Thus the clustering algorithm helps especially for those maps that are most difficult to compress.

5. Variants

5.1 Limited number of trees with depth ≤ 2

We now consider the possibility of other time/space tradeoffs, intermediate between using no clustering (fast decompression but needing more space) and clustering with the MST (reducing storage requirements, but with slower decompression). One of the problems with using the forest structure of Section 3 is that the number of clusters may be fairly large. In fact, there are many maps which don't benefit from

XORing with some other map and these form singleton-clusters. In a practical application, we may want to limit the number of clusters, so that the cluster index can be represented by just a few bits. Another problem with the MST approach is that some of the trees produced may be deep, as can be seen from the examples of the previous section; for such trees, the decompression time for a given bitmap is proportional to its level in the tree. Thus it would be desirable to create trees similar to the ones discussed above, but restricted in number and constrained to have a small depth; we shall examine the case in which the maximum depth is two. The corresponding forest is not necessarily optimal, so the space complexity will probably increase; on the other hand, at most one additional map is needed for decompressing any map, so we gain in speed.

This new variant is described more precisely as follows. Suppose we can partition the file \mathcal{B} of m bitmaps into n mutually disjoint clusters B_1, \dots, B_n . In each cluster one map $b_i \in B_i$ is chosen as a *representative*; thus we have a mapping $\Pi : \mathcal{B} \rightarrow \{(B_1, b_1), \dots, (B_n, b_n)\}$, where $b_i \in B_i$ and $\{B_i\}$ is a partition of \mathcal{B} . In terms of the preceding discussion, each B_i corresponds to a tree of depth at most two, and b_i is the root of the tree B_i . For each B_i the associated b_i will be compressed directly and stored. If $x \neq b_i$ is a bitmap in B_i , it is first XORed with the representative b_i , and then the result is compressed and stored, together with the index i of the cluster to which it belongs.

The problem, then, is to find that partition of \mathcal{B} into clusters, and the selection of the designated representatives, for which the savings are greatest. We have not found an efficient optimal solution to the problem and believe it to be very hard. To support this assertion, we will approximate this problem by a similar one, and show that the second problem is hard. We expect that good partitions will have certain reasonable properties. One way to identify a good partition is by searching for a partition that satisfies such a set of properties. We will now define our second problem as an illustration of this approach.

We again use the Hamming distance between maps x and y , $\text{HD}(x, y)$, as our measure of dissimilarity. Our revised problem is to find a partition of \mathcal{B} into n sets and to select n representatives, such that these properties are satisfied: (1) the total number of 1-bits in the representatives, $\sum_{i=1}^m \text{HD}(b_i, X_0)$, is minimized; (2) the sum of all the distances from a map to the corresponding representative of the cluster it belongs to, $\sum_{i=1}^n \sum_{V \in B_i} \text{HD}(V, b_i)$, is minimized; and (3) clusters should also be “different” enough, in the sense that the distance between representatives of different clusters should be bounded below by some constant.

At first sight, this restricted problem seems easier than the previous one, where, since neither the number, nor the depth of the trees were restricted, we must locate the optimal member from a much larger set. Once the partition of \mathcal{B} into n clusters is given, we now only have to check every possible choice of a representative, instead of generating all the spanning trees in each cluster. We show in the Appendix, however, that the restricted problem of finding the optimal partition is NP-complete, which

implies that there is probably no feasible (polynomial) algorithm available.

Instead of the optimizing problem, we actually define the following corresponding decision problem.

HDC (Hamming Distance Clustering): *Instance:* A set \mathcal{B} of m bitmaps of length ℓ bits each, and integers K_1, K_2, K_3 and n . *Question:* Is there a partition of \mathcal{B} into n disjoint subsets B_1, \dots, B_n , such that in each of them there is one representative $b_i \in B_i$ satisfying the following conditions:

- (1) the representatives are sparse: $\sum_{i=1}^n \text{HD}(b_i, 0) \leq K_1$;
- (2) the total distance is bounded: $\sum_{i=1}^n \sum_{V \in B_i} \text{HD}(V, b_i) \leq K_2$;
- (3) cluster representatives are not close: $\forall i \neq j \quad \text{HD}(b_i, b_j) \geq K_3$;
- (4) in terms of the given representatives, each map is in the best possible cluster:

$$\forall V \in \mathcal{B} \quad \forall i (V \in B_i \rightarrow \text{HD}(V, b_i) = \min_{j \in \{1, \dots, n\}} \text{HD}(V, b_j)) ?$$

Theorem 1. *HDC is NP-complete.*

Proof: Given the partition of \mathcal{B} into clusters, condition (1) can be checked in time $O(n)$, conditions (2) and (4) in time $O(m\ell)$ and condition (3) in time $O(n^2\ell)$. Thus $\text{HDC} \in \text{NP}$. For the reduction, we use the following problem known to be NP-complete (see [8]):

3DM (3 Dimensional Matching): *Instance:* A set $M \subseteq W \times X \times Y$, where W, X and Y are disjoint sets having the same number q of elements. *Question:* Does M contain a *matching*, that is, a subset $M' \subseteq M$ such that $|M'| = q$ and no two elements of M' agree in any coordinate?

We show that $3\text{DM} \propto \text{HDC}$. Consider an arbitrary instance of 3DM. We construct an instance of HDC. Create a set of m bitmaps, \mathcal{B} , as follows: there is one bitmap for each element in M (thus $m = |M|$); each bitmap has $\ell = 3q$ positions, the first q positions corresponding to the elements of W , the next q to the elements of X and the last q to the elements of Y . In the bitmap corresponding to the element $(w, x, y) \in M$, where $w \in W, x \in X$ and $y \in Y$, there are 1's in the positions corresponding to w, x and y , and zeros elsewhere; hence each bitmap consists of exactly three 1's and $3(q-1)$ zeros. Set $K_1 = 3q, K_2 = 4(m-q), K_3 = 6$ and $n = q$.

Suppose there exists a matching M' . Denote the q rows corresponding to the matching by b_1, \dots, b_q : these will be the representatives of the q clusters B_1, \dots, B_q respectively. Now define the partition of \mathcal{B} in the following way. If a map $V \in \mathcal{B} - \bigcup_{i=1}^q b_i$ has two overlapping 1's with one of the representatives, say b_j , then V is included in B_j ; such a b_j is unique. Otherwise, choose the (unique) map $b_v, v \in \{1, \dots, q\}$, which has its leftmost 1 in the same bit-position as V , and include V in the cluster B_v . (1) Since each bitmap contains exactly three 1-bits, this is in particular

true for b_1, \dots, b_q . Thus $\sum_{i=1}^q \text{HD}(b_i, 0) = 3q = K_1$. (2) As to the remaining $m - q$ maps in $\mathcal{B} - \bigcup_{i=1}^q b_i$, each such map V has either one or two 1's overlapping with the representative b_v of the cluster to which V belongs, so that $\text{HD}(V, b_v)$ is either 4 or 2 respectively. Thus $\sum_{i=1}^q \sum_{V \in B_i} \text{HD}(V, b_i) \leq 4(m - q) = K_2$. (3) Since for $i \neq j$, b_i and b_j correspond to different elements of the matching M' , they have no overlapping 1's, thus $\text{HD}(b_i, b_j) = 6 = K_3$ for all $i \neq j$. Condition (4) follows directly from the construction and the obvious fact that no map can have more than two overlapping 1's with any other map.

Conversely, suppose that there is a partition of \mathcal{B} into n clusters, with representatives b_1, \dots, b_n that satisfy (1), (2), (3) and (4). Define M' as the elements of M corresponding to the bitmaps b_1, \dots, b_n . Since for $i \neq j$, $\text{HD}(b_i, b_j) = 6$ and there are only three 1-bits in each map, b_i and b_j cannot have 1-bits in identical positions. Thus the corresponding elements of M' do not agree in any coordinate, and since there are $n = q$ such elements, it follows that the set M' is a matching. ■

5.2 Artificial representatives

Suppose we can find a partition of the set of bitmaps into clusters of similar bitmaps. Instead of choosing one element in each cluster as a representative which is processed differently from the other maps in the cluster, we could construct a new, artificial bitmap R_i for each cluster, which will act as that cluster's representative. Thus R_i can be adapted to optimally fit the entire cluster. All the maps in the cluster are now processed in the same way. As in the original problem, the quantity we try to minimize is the number of 1-bits in the maps obtained by XORing the original maps with their respective R_i , plus the number of 1-bits in the representative R_i .

Given a set of clusters, it is easy to find the optimal R : we proceed by adjoining a bitmap B_0 , containing only zeros, to each cluster. Our problem is then, given the bitmaps B_0, B_1, \dots, B_k in a cluster, to create a vector R such that $\sum_{i=0}^k \text{HD}(B_i, R)$ is minimized; B_0 is included so the ones in R are included in the sum. The minimization can be carried out individually for each bit position. Let $B_i = (b_1^i, b_2^i, \dots, b_\ell^i)$ and $R = (r_1, \dots, r_\ell)$. For each $j = 1, \dots, \ell$, we are looking for r_j such that $\sum_{i=0}^k |b_j^i - r_j|$ is minimized.

The optimal r_j is obtained from the original b_j^i ($1 \leq i \leq k$) by a simple majority vote. That is, if more than half of the bits in the j -th position of the maps B_i are ones, then r_j should be one. If there is a tie, or if the majority of values is zero, then r_j should be zero. The value of r_j can be represented compactly as $\lfloor 2 \sum_{i=0}^k b_j^i / (m + 1) \rfloor$. (This is a special case of the well known result that the median of a set of numbers is the value that minimizes the sum of the absolute derivations from that value.)

We just saw that given a partition into clusters, it is easy to construct a set of representatives by using the majority rule. The converse problem would be, given a set of representatives, to find an optimal partition into clusters. This is easily done by

checking, for each bitmap, to which representative it is closest. However, if one must determine both the optimal representatives, and their clusters, with restrictions given on the number of 1-bits in the representatives, then one again has an NP-complete problem.

PCFD (Partition into Clusters of Fixed Density): *Instance:* A set $\mathcal{B} = \{b_1, \dots, b_m\}$ of bitmaps of length ℓ bits each, n integers d_0, \dots, d_{n-1} , and an integer K . *Question:* Can \mathcal{B} be partitioned into n clusters B_0, \dots, B_{n-1} such that there exists a representative r_i for cluster B_i with exactly d_i 1-bits and such that the total Hamming distance is bounded, $\sum_{i=0}^{n-1} \sum_{b \in B_i} \text{HD}(b, r_i) \leq K$?

Theorem 2. *PCFD is NP-complete.*

Proof: Once the partition and the representatives are given, the condition is easily checked, thus PCFD \in NP. For the reduction we use the *Partition* problem [8].

PAR (Partition): *Instance:* A set of positive integers $\mathcal{A} = \{a_1, \dots, a_k\}$. *Question:* Is there a subset J of $I = \{1, \dots, k\}$, such that $\sum_{i \in J} a_i = \sum_{i \in I-J} a_i$?

Let $A = \sum_{i=1}^k a_i$. If A is odd, the partition problem is trivial, so we assume that A is even.

We show that PAR \propto PCFD. Letting a_1, \dots, a_k be an instance of PAR, we construct an instance of PCFD. We set $m = k$ and define the set \mathcal{B} of m binary vectors b_1, \dots, b_m , each of length $\ell = A$, in the following way:

$$b_i = (0, 0, \dots, 0, 1, 1, \dots, 1, 0, 0, \dots, 0)$$

$$\underbrace{\hspace{10em}}_{\sum_{j=1}^{i-1} a_j \text{ times}} \quad \underbrace{\hspace{10em}}_{a_i \text{ times}} \quad \underbrace{\hspace{10em}}_{\sum_{j=i+1}^m a_j \text{ times}}$$

That is, the vector b_i consists, in order, of $\sum_{j=1}^{i-1} a_j$ zeros, followed by a_i 1's, followed by $\sum_{j=i+1}^m a_j$ zeros, where we use the convention that $\sum_{i=x}^y a_i = 0$ if $x > y$. Let $n = 2$, $d_0 = d_1 = A/2$, and $K = \left(\frac{m}{2} - 1\right) A$.

Suppose there exists a set J as required. Let δ_{iJ} be defined by

$$\delta_{iJ} = \begin{cases} 1 & \text{if } i \in J \\ 0 & \text{otherwise,} \end{cases}$$

and construct the vector r_0 to consist, in order, of a_1 occurrences of δ_{1J} , followed by a_2 occurrences of δ_{2J} , \dots , followed by a_m occurrences of δ_{mJ} . Define r_1 as the binary complement of r_0 . The number of 1's in r_0 is thus $\sum_{i \in J} a_i = A/2 = d_0$, and therefore the number of 1's in r_1 is also $A/2 = d_1$. b_i is in the cluster associated with r_0 if

$a_i \in J$; else, b_i is in the cluster represented by r_1 . Summing the Hamming distances, we get

$$\sum_{i \in J} \text{HD}(b_i, r_0) + \sum_{i \in I-J} \text{HD}(b_i, r_1) = \sum_{i=1}^m \left(\frac{A}{2} - a_i \right) = m \frac{A}{2} - A = K,$$

which shows that \mathcal{B} can be partitioned as required in the definition of PCFD.

Conversely, suppose the maps b_i can be partitioned into two clusters and that there exist two representatives, r_0 and r_1 , each with $A/2$ 1-bits, such that the total Hamming distance is bounded by K . Since r_0 has $A/2$ 1-bits, the Hamming distance of b_i to r_0 is at least $A/2 - a_i$. This minimum can be achieved only if all the 1-bits in b_i are in positions where r_0 has also 1-bits. The same remark is true for r_1 . On the other hand, if even one of the b_i has a Hamming distance greater than $A/2 - a_i$ to the representative of the cluster it belongs to, the total Hamming distance would be larger than $\sum_{i=1}^m (A/2 - a_i) = K$.

It follows that r_0 and r_1 have 1-bits in all the positions where the maps belonging to the corresponding clusters have 1-bits. On the other hand, neither r_0 nor r_1 can have 1-bits in other positions, because the positions of the 1-bits in all the maps together cover all the A positions and there are only A 1-bits in $r_0 \cup r_1$. Thus r_1 must be the binary complement of r_0 , and the 1-bits in both r_0 and r_1 are arranged in blocks corresponding to the bit-positions covered by the vectors b_i . This can be reformulated as

$$\forall i \in I \quad \forall j \in \{0, 1\} \quad \text{HD}(b_i, r_j) = \frac{A}{2} - a_i \iff \text{HD}(b_i, r_{(1-j)}) = \frac{A}{2} + a_i.$$

We therefore can define the set $J \subseteq I$ to contain the indices of the maps b_i which have overlapping 1-bits with one of the representatives, say, r_0 , i.e.,

$$i \in J \iff \text{HD}(b_i, r_0) = \frac{A}{2} - a_i.$$

The positions of the 1-bits in different b_i are mutually disjoint, and the positions of the 1-bits of b_i , with $i \in J$, are included in the set of the positions of the 1-bits of r_0 . Therefore the number of 1-bits in $\bigcup_{i \in J} b_i$, $\sum_{i \in J} a_i$, is equal to the number of 1-bits in r_0 , $A/2$. But then $\sum_{i \in I-J} a_i$ is equal to $A - A/2 = A/2$. Hence

$$\sum_{i \in J} a_i = \sum_{i \in I-J} a_i,$$

which shows that the set \mathcal{A} can be partitioned as requested. ■

6. Conclusion and Future Work

We have presented a new algorithm for transforming a set of bitmaps, which in principle may improve any previous compression method that does not take into account possible interrelationships among the different bitmaps. The experimental results suggest that the new method is particularly effective for bitmaps which are not extremely sparse. This may have several applications.

For example, bit-slices of signature methods are often chosen so that the density of 1-bits is $\frac{1}{2}$ [15]. Such vectors are almost impossible to compress individually. There may however be a possible gain by using clustering. Also, the possibility of compression would permit us to increase the size of the signature, resulting in more efficient retrieval, without affecting the space requirements [3]. Another application would be to IR bitmaps which have already been slightly compacted, such as the maps obtained by applying one iteration of the hierarchical bit-vector compression technique referred to in Section 2.1. Finally, there might be applications to areas outside of IR, such as image compression, where adjacent raster rows may be similar, or processing of genetic information, where different DNA strings often share long identical substrings.

We noted several situations in which it would be desirable to put constraints on the graph structure produced by our algorithm. However, we concluded that the restricted problem is most probably intractable in reasonable time, thus it is justified to search for feasible heuristics. The possibility of using such heuristics will be investigated in future research.

References

- [1] **Abramowitz M., Stegun I.A.**, *Handbook of Mathematical Functions*, Dover Publishing, Inc., New York (1965).
- [2] **Bell T., Witten I.H., Cleary J.G.**, Modeling for Text Compression, *ACM Computing Surveys* **21** (1989) 557–591.
- [3] **Bookstein A., Klein S.T.**, Using Bitmaps for Medium Sized Information Retrieval Systems, to appear in *Inf. Proc. and Management* (1990).
- [4] **Choueka Y., Fraenkel A.S., Klein S.T., Segal E.**, Improved hierarchical bit-vector compression in document retrieval systems, *Proc. 9-th ACM-SIGIR Conf.*, Pisa; ACM, Baltimore, MD (1986) 88–97.
- [5] **Choueka Y., Fraenkel A.S., Klein S.T., Segal E.**, Improved Techniques for Processing Queries in Full-Text Systems, *Proc. 10-th ACM-SIGIR Conf.*, New Orleans (1987) 306–315.
- [6] **Faloutsos C., Christodoulakis S.**, Signature files: An access method for documents and its analytical performance evaluation, *ACM Trans. on Office Inf.*

- Systems* **2** (1984) 267–288.
- [7] **Fraenkel A.S., Klein S.T.**, Novel Compression of sparse Bit-Strings, in *Combinatorial Algorithms on Words*, NATO ASI Series Vol **F12**, Springer Verlag, Berlin (1985) 169–183.
 - [8] **Garey M.R., Johnson D.S.**, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W.H. Freeman, San Francisco (1979).
 - [9] **Jakobsson M.**, Huffman coding in Bit-Vector Compression, *Inf. Processing Letters* **7** (1978) 304–307.
 - [10] **Klein S.T., Bookstein A., Deerwester S.**, Storing Text Retrieval Systems on CD-ROM: Compression and Encryption Considerations, *ACM Trans. on Information Systems* **7** (1989), 230–245.
 - [11] **Knuth D.E.**, *The Art of Computer Programming, Vol I, Fundamental algorithms*, Addison-Wesley, Reading, Mass. (1973).
 - [12] **Kruskal J.B.**, On the shortest spanning subtree of a graph and the Travelling Salesman Problem, *Proc. Amer. Math. Soc.* **7** (1956) 48–50.
 - [13] **Lelewer D.A., Hirschberg D.S.**, Data Compression, *ACM Computing Surveys* **19** (1987) 261–296.
 - [14] **Schuegraf E.J.**, Compression of large inverted files with hyperbolic term distribution, *Inf. Proc. and Management* **12** (1976) 377–384.
 - [15] **Stiassny S.**, Mathematical analysis of various superimposed coding methods, *Amer. Documentation* **11** (1960) 155–169.
 - [16] **Storer J.A.**, *Data Compression: Methods and Theory*, Computer Science Press, Rockville, Maryland (1988).
 - [17] **Teuhola J.**, A Compression method for Clustered Bit-Vectors, *Inf. Processing Letters* **7** (1978) 308–311.
 - [18] **Vallarino O.**, On the use of bit-maps for multiple key retrieval, *SIGPLAN Notices, Special Issue* Vol. **II** (1976) 108–114.
 - [19] **Wedekind H., Härder T.**, *Datenbanksysteme II*, B.-I. Wissenschaftsverlag, Mannheim (1976).
 - [20] **Yao A.C.C.**, An $O(|E| \log \log |V|)$ algorithm for finding minimum spanning trees, *Inf. Processing Letters* **4** (1975) 21–23.