



HAL
open science

Compressing bipartite graphs with a dual reordering scheme

Maximilien Danisch, Ioannis Panagiotas, Lionel Tabourier

► **To cite this version:**

Maximilien Danisch, Ioannis Panagiotas, Lionel Tabourier. Compressing bipartite graphs with a dual reordering scheme. 2022. hal-03875560

HAL Id: hal-03875560

<https://hal.science/hal-03875560>

Preprint submitted on 28 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Compressing bipartite graphs with a dual reordering scheme

Maximilien Danisch, Ioannis Panagiotas, Lionel Tabourier*

Sorbonne Université, CNRS, LIP6, F-75005 Paris, France

Abstract

In order to manage massive graphs in practice, it is often necessary to resort to graph compression, which aims at reducing the memory used when storing and processing the graph. Efficient compression methods have been proposed in the literature, especially for web graphs. In most cases, they are combined with a vertex reordering pre-processing step which significantly improves the compression rate. However, these techniques are not as efficient when considering other kinds of graphs. In this paper, we focus on the class of bipartite graphs and adapt the vertex reordering phase to their specific structure by proposing a dual reordering scheme. By reordering each group of vertices in the purpose of minimizing a specific score, we show that we can reach better compression rates. We also suggest that this approach can be further refined to make the node orderings more adapted to the compression phase that follows the ordering phase.

1 Introduction

Many real-world systems are adequately represented by graphs, as they allow to model interacting entities. Among many examples, graph representations are used to describe the world wide web as webpages connected by hyperlinks, social networks where accounts are connected by a friendship or a follower relationship, human activity on online platforms, words occurring in a same sentence, or even biochemical reactions between proteins. As data is generated faster and in larger quantities than ever before, we are led to handle massively increasing graph sizes. For example, graphs from social networks such as Facebook or Twitter demand up to terabytes of RAM to be loaded in main memory. It is impossible to process them with standard computers, making the retrieval and analysis of relevant information more challenging. In this context, designing methods to

*Authors contribution: MD proposed the original idea of the paper, IP realized the codes and experiments, IP and LT contributed equally to the analysis and the writing of the paper.

produce compact representations of the information contained in a graph has emerged as an important research question, named *graph compression*. The compression can be either *lossy* or *lossless*, depending on whether part of the information is left out of the representation or not. Here, we are interested in lossless compression.

Very efficient methods for lossless graph compression have been proposed in the literature [6, 29]. These methods have been designed primarily to be applied to web graphs, *i.e.* graphs where nodes are URLs and directed edges represent hyperlinks. In this case, nodes are naturally ordered according to the lexicographic order of the URLs. It was observed that this natural vertex ordering favors good compression rates notably because it satisfies two properties that compression techniques exploit: *locality* and *similarity* [6]. The first term means that a vertex is mainly connected to vertices with a close index, while the second means that vertices which share many neighbors have close indexes. Unfortunately, other types of networks do not necessarily exhibit such an adequate natural vertex ordering, leading to much less efficient compressed representations. That is why vertex reordering techniques have emerged as an essential preprocessing phase to achieve high graph compression rates [10, 5, 13]. These techniques seek to find an ordering that satisfies as much as possible the properties of locality and similarity, as well as any other properties which could benefit the subsequent compression phase.

In this work, we focus on the problem of lossless compression of bipartite graphs and especially on the reordering phase of the process. Bipartite graphs are appropriate models to represent systems where relations connect two different kinds of entities: users consuming contents on an online platform, individuals and groups to which they are affiliated, indexes referencing pages, etc. We utilize the fact that both groups of nodes play very different roles in the graph and propose reordering the two groups independently using two distinct objectives, rather than optimizing a single objective as is usual. These objectives are defined to be consistent with the logic of the compression methods that follows.

The rest of the paper is organized as follows. In Section 2 we provide the background and summarize related work. Then in Section 3, we present the dual reordering scheme that we propose to compress bipartite graphs efficiently. Section 4 is dedicated to the description of the experimental protocol and results that we have obtained, which show that the scheme is indeed more efficient than state-of-the-art reordering methods for this purpose. We present in Section 5 some leads to develop this reordering method further before briefly concluding in Section 6.

2 Background and related work

Preliminaries. We first define the basic vocabulary needed to describe the problem of graph compression. Let $G = (V, E)$ be an undirected graph where V is the set of vertices and E the set of edges. We call $u, v \in V$ neighbors in G if the edge (u, v) exists in E . The *neighborhood* of a vertex u , denoted $\mathcal{N}(u)$, is the

set of its neighbors and the degree of u , d_u , is the size of this neighborhood. A graph is *bipartite* if V can be partitioned into two disjoint sets V_\top and V_\perp such that $V = V_\top \cup V_\perp$ and every edge in E contains one vertex from V_\top and one from V_\perp . V_\top and V_\perp are called respectively *top* and *bottom* nodes. As previously mentioned, graph compression is tightly linked to the node indexation of the graph. We define an ordering of V as an injective function $\pi : V \rightarrow \{1, \dots, |V|\}$, *i.e.*, a renumbering of the vertices of V . Vertex u precedes v if $\pi(u) < \pi(v)$ and follows v when $\pi(u) > \pi(v)$.

Note that graphs built from real-world data are known to be sparse. A straightforward way to store efficiently sparse graphs in memory is to represent them as adjacency lists *i.e.*, the lists of neighbors of each node in the graph. This format is the starting point of efficient lossless graph compression methods, which then use a combination of techniques, that we discuss below, to improve upon it.

Compression methods for web graphs. In the case of web graphs, Boldi and Vigna’s *WebGraph* framework [6] identified several lossless compression techniques that have allowed to reach high compression rates, which are still competitive with current state-of-the-art methods. Indeed, the compressed representation of such graphs requires as few as 2–3 bits per edge¹. The general idea underlying this framework is to take advantage of two central characteristics of web graphs. The first one, named *locality*, supposes that connected vertices have relatively close indexes. In lexicographically ordered web graphs, it is usually the case as the source and target of a hyperlink are often part of a same domain name. The second one, *similarity*, means that similar vertices (nodes which share a large subset of their neighbors) appear close to each other in the ordering. This is also the case in web graphs because two pages of a same domain tend to have close index according to the lexicographic order and these two pages commonly have very similar navigational links.

Among the compression techniques developed for WebGraph, Boldi and Vigna defined the *zeta encoding*, a universal code² especially designed for graphs which encodes a sequence of consecutive integers as intervals (*i.e.*, represented only with two values). They have also incorporated a widely-used compression technique called *delta encoding* [10, 29]: when storing a list of sorted integers u_1, \dots, u_k , delta encoding consists in representing it as a list of consecutive gaps *i.e.*, $u_1, (u_2 - u_1 - 1), \dots, (u_k - u_{k-1} - 1)$. Because these gaps are likely to be much smaller than the integers of the original list, storing them with variable-length quantity encoding schemes usually requires less space. This is particularly efficient if the neighbors of a node have close index, which is consistent with the locality principle. Another crucial technique for graph compression – especially in the perspective of our work – is *referencing*. Referencing allows a vertex u to select one of its predecessors v within a fixed window size as its reference, and

¹Note that lossless compression is usually empirically measured in bits per edge, as the total size of the graph is not necessarily easy to interpret.

²A universal code is a prefix code such that the expected lengths of the words are within a constant factor of the expected lengths of the words using the optimal code.

store their common neighbors implicitly. Practically, one keeps for the node u a 0 – 1 bit array b of size d_v where $b[i]$ is 1 if u and v share the i th neighbor of v and 0 otherwise. The neighbors of u which are not neighbors of its reference v (called residual neighbors) are stored explicitly after the bit array b . When u and v have many common neighbors, such a representation is more efficient than explicitly storing all of u 's neighbors.

A few other approaches can be mentioned regarding web graph compression. On the one hand, some of them develop on Boldi and Vigna's framework. For instance, Grabowski and Bieniecki [17] handle referencing in an alternative way: a vertex u can only reference its immediate predecessor v , but each copied neighbor of v is allocated more than one bit, so as to portray more information about the neighborhood of u . Also, Liakos *et al.* [22] store the denser diagonal part of the adjacency matrix of the graph separately for efficiency purposes and resort to using the WebGraph framework for the rest of the graph. Recently, Versari *et al.* [29] proposed Zuckerli, a new software for compression that incorporates several improvements to the WebGraph framework. These improvements notably consist in employing a hybrid encoding scheme for storing integers and improving the referencing algorithm. They have shown empirical improvements ranging from 10% to 20% on many instances over WebGraph and other methods. On the other hand, some works are based on schemes which are radically different from WebGraph, such as Brisaboa *et al.* [7] which introduces k^2 -trees to compress web graphs. Their idea is to recursively partition the adjacency matrix of a graph into k^2 parts and stop when a submatrix consists solely of zeros or ones. By storing the resulting tree in an adequate structure [12], they managed to achieve efficient compression rates. Another approach is using *virtual nodes* to represent frequently appearing structures in the graph. For instance, Buehrer and Chellapilla [9] replace bicliques with such virtual nodes to encode them more efficiently, then the remaining edges are encoded using standard web graph encoding scheme, such as the ones mentioned earlier, whereas Rossi and Zhou [28] represent cliques compactly. Claude and Navarro [11] take inspiration from grammar-based compression to adapt the Re-Pair [21] algorithm to graphs: Re-Pair continuously replaces a pair of symbols (here, vertices) that appear together frequently in adjacency lists with a new symbol. Aside from the method discussed above, Grabowski and Bieniecki [17] proposed another compression scheme, which separates vertices into blocks and "merges" the neighborhoods of the vertices in each block so as to remove duplicate information.

A more extensive coverage of various approaches to lossless compression on web graphs can be found in a survey by Besta and Hoefler [3].

Inverted index compression. Among other families of data that have attracted much interest for lossless compression purposes, inverted indexes are essential data structures for the efficient implementation of many information retrieval tasks. They are used to index large collections of documents in the form of an ordered list of IDs, corresponding to the documents where a specific word appear. While this description stresses the asymmetry between the words

(that we call *queries* for generality) and the documents (*data*), an inverted index can be described and stored in the same way as a bipartite graph, where a node either represents a piece of data or a query. Note that inverted indexes may be enriched with additional information, but we are not interested in this aspect here as we focus on what the bipartite structure brings.

According to a recent review [25], the inverted index compression process can be split into three main parts: i) compressing a single integer, ii) compressing a list of integers and iii) compressing many lists together. These steps can be mapped to the techniques described in Boldi and Vigna WebGraph framework. i) The combination of delta and zeta encodings aim at compressing single integers. This combination is particularly efficient because zeta-encoding is well-suited for small integers, while delta encoding ensures to have small integers if the graph has high locality. Yet other integer encoding methods are available such as Elias gamma and delta codes [15] or Rice code [27] and their offsprings. ii) Compressing a list of integers usually relies on clustering consecutive integers to use a summarized representation of the cluster. For that purpose, WebGraph uses notably a variation of binary packing, because their coding method implies the existence of long sequences of 0 and 1. Here again, other methods are possible, such as entropy coding techniques, among which Huffman coding [23] or Arithmetic coding [24] are probably the most famous. iii) The referencing approach used in WebGraph is one of the usual ways to compress an ensemble of similar enough lists. Among other possibilities, we point to Pibiri *et al.* [26] who proposed a dictionary based approach that yields good results for this part of the process.

In short, we observe that the principle of the methods developed for web graphs are similar to those used for inverted index compression. Inverted indexes can be seen as bipartite graphs, which makes these techniques relevant for our problem. Moreover, in the case that we consider, we will have the liberty to use nodes indifferently as queries or data.

Vertex reordering. Other types of networks have specific structural properties which can be used to achieve efficient compression. For instance, in the case of social networks, Chierichetti *et al.* [10] exploit the fact that these networks often exhibit high *reciprocity* [4], meaning that if there is a directed edge from node i to node j , there is a high probability that there is also one from j to i . So it is possible to improve the compression by simply signaling that a link is reciprocal and discarding the explicit reciprocal link. Importantly, the authors of [10] also notice that the methods described for web graphs usually works well because the data collection ordering respects to some extent similarity and locality. Unfortunately, the collection ordering with other data – and particularly social data – depends on the crawling procedure and it is not guaranteed to be as efficient in terms of locality and similarity. Consequently, they suggest that before applying the usual compression techniques, one should reorder the vertices of the graph to restore as much as possible similarity and locality.

An efficient delta encoding supposes that sibling nodes – i.e. nodes which are

neighbors of a same node – have close indexes, reciprocally, good referencing demands that nodes with close indexes share many neighbors. The similarity principle can ensure these two properties. These observations highlight the fact that vertex ordering plays an important role in graph compression. Besides, as the vertex reordering phase is done only once in general, the question of additional reordering time cost is usually irrelevant since it can be done offline. This contrasts with the compression and decompression times which are independent of the vertex ordering method. In this spirit, Apostolico and Drovandi [2] proposed GBFS which reorders the vertices using breadth-first search in a preliminary step to improve compression. Also, Boldi *et al.* [5] suggested an approach called LLP inspired by community detection algorithms, which partitions vertices into communities and puts vertices of a same community in consecutive positions in π . Several other reordering methods relying on a similar community-based strategy are discussed in [3].

These techniques aim at optimizing the quality of the vertex ordering in the perspective of future processings, the first of which being delta encoding. It is therefore natural to express this issue as an optimization problem on orderings. Chierichetti *et al.* [10] introduced the MINLOGGAPA problem: given a graph G , MINLOGGAPA seeks the ordering π that minimizes the objective function:

$$\sum_{v \in V} f_{\pi}(v, \mathcal{N}(v)) = \sum_{v \in V} \sum_{i=1}^{d_v-1} \log(\pi(u_{i+1}) - \pi(u_i)).$$

where vertices u_1 to u_{d_v} are the neighbors of vertex v sorted by increasing value of π such that any consecutive difference is positive. By minimizing the MINLOGGAPA objective function, we minimize the gaps between neighbors, and it is expected that the delta encoding described above will be more efficient.

In [10], the authors propose a heuristic called *shingle*, which computes a fingerprint of the neighborhood per node and re-orders nodes with similar fingerprints near each other. We describe it with some more details as we use its principles later in this work. Supposing that σ is a random permutation of a set, then the smallest element in the set according to σ is defined as its *shingle*. It has been shown [8] that the Jaccard coefficient $J(A, B) = |A \cap B|/|A \cup B|$ is the probability that the sets A and B have the same shingle. So a shingle – or several shingles generated with different permutations, depending on the desired precision – can be used to fingerprint the set of neighbors of the nodes in the purpose of evaluating approximate Jaccard index. As the Jaccard index is a measurement of the similarity of two sets, it is used to evaluate the overlap between the neighborhoods of pairs of nodes. In other words, a shingle is used to approximate $J(\mathcal{N}(u), \mathcal{N}(v)) = |\mathcal{N}(u) \cap \mathcal{N}(v)|/|\mathcal{N}(u) \cup \mathcal{N}(v)|$, which in turn measures the overlap between the neighborhoods of nodes u and v . The shingle re-ordering technique consists in ordering nodes in such a way that consecutive nodes have a high Jaccard coefficient, thus leading to relatively high level of similarity between nodes. Note that in practice, Chierichetti *et al.* [10] use a combination of hash functions instead of random permutations to approximate Jaccard coefficients. It is a fast reordering method and yields an efficient mini-

mization of the objective function of the MINLOGGAPA problem in comparison to the natural ordering (*i.e.*, the vertices order as produced by the data collection process). They suggested, but did not prove, that the MINLOGGAPA problem is NP-hard. It was later shown by Dhulipala *et al.* [13] and they also proposed an efficient recursive bisection heuristic to tackle it practically, that we develop in more details below as we make use of it in this paper.

The case of recursive bisection reordering. Dhulipala *et al.* [13] method targets compression of undirected graphs in general, however it is based on optimizing an objective for inverted indexes. Indeed, they first transform the original graph $G_u = (V, E)$ into an inverted index; it is essentially a bipartite graph $G = (\mathcal{Q} \cup \mathcal{D}, E')$, where its vertices are either representing queries ($q \in \mathcal{Q}$) or data ($u \in \mathcal{D}$). Then, they adapt the MINLOGGAPA optimization problem to this context, by formulating the BIMLOGGAPA problem, which seeks an ordering of the nodes in \mathcal{D} to minimize

$$\sum_{q \in \mathcal{Q}} f_\pi(q, \mathcal{N}(q)) = \sum_{q \in \mathcal{Q}} \sum_{i=1}^{d_q-1} \log(\pi(u_i) - \pi(u_{i+1})). \quad (1)$$

They propose a recursive bisection heuristic denoted RECBIS to tackle this minimization problem. Their heuristic adapts the well-known Kernighan-Lin [19] and Fidduchia-Mattheyses [16] heuristics for graph partitioning to the problem of vertex reordering. In a few words, Dhulipala *et al.*'s heuristic starts from an initial partition of the set \mathcal{D} into two equal-sized sets \mathcal{D}_1 and \mathcal{D}_2 (so a random balanced assignment). Assuming \mathcal{D} 's vertices are to be ordered between π_l and π_r in the final ordering π , vertices of \mathcal{D}_1 are in positions $\pi_l, \dots, \pi_l + |\mathcal{D}_1| - 1$ and vertices of \mathcal{D}_2 are in positions $\pi_l + |\mathcal{D}_1|, \dots, \pi_r$. The heuristic consists in exchanging vertices between the two partitions as long as they diminish the cost of the objective function. Note that RECBIS does not directly optimize equation (1), but aims to optimize the expected cost instead. It is equivalent to considering that the neighbors of q are arranged regularly at equal distance from each other. This is a simplification for the fact that the exact location of these vertices inside \mathcal{D}_1 and \mathcal{D}_2 are not yet known at that step. Once the partition of \mathcal{D} into \mathcal{D}_1 and \mathcal{D}_2 has been determined, the process is repeated recursively on the two inverted indexes $G_1 = (\mathcal{Q} \cup \mathcal{D}_1, E_1)$ and $G_2 = (\mathcal{Q} \cup \mathcal{D}_2, E_2)$. The internal ordering of the vertices of \mathcal{D}_1 and \mathcal{D}_2 is therefore fixed in the subsequent recursion calls.

In [13], the authors also mention using recursive bisection to compress actual bipartite graphs, but there is no detail about the specifics of the process on such graphs. The BIMLOGGAPA problem focuses mostly on ordering properly vertices of the data group \mathcal{D} , which is relevant when compressing inverted indexes. Our approach consists in ordering both the top (V_\top) and bottom (V_\perp) sets of nodes, so that if we represent the bipartite graph with adjacency lists of nodes in V_\top , ordering V_\top nodes improves the compression by benefiting from delta encoding, while ordering nodes in V_\perp improves the compression by benefiting from referencing.

3 Dual reordering scheme

In this section we describe our approach to vertex reordering in bipartite graphs. As suggested in Section 2, the reordering process in the case of bipartite graphs aims at reorganizing nodes in a way that would improve the property of similarity³: we want vertices with close indexes to share many common neighbors.

We may choose to represent the graph either by storing adjacency lists of \perp nodes or those of \top nodes. Note that our goal is to improve compression regardless of what the vertices represent, so one may try both V_{\perp} and V_{\top} as the query set, and then keep the option yielding the best compression rate. Without loss of generality, we assume from now on that V_{\perp} is the set of queries and that V_{\top} is the set of data entries. That is, for each node $u \in V_{\perp}$, we store the list of its neighbors $\mathcal{N}(u)$ in V_{\top} . Under this representation, all edges in the graph are stored exactly once. We set that π is the permutation ordering V_{\top} nodes and ϕ the permutation ordering V_{\perp} nodes.

First, we describe how we design the π ordering on V_{\top} nodes, which is essentially based on RECBIS with a few improvements. Then, our efforts focus on finding ϕ *i.e.*, the ordering of V_{\perp} , with the purpose of maximizing the impact of referencing.

3.1 Top nodes ordering to improve delta encoding

\top nodes (*i.e.*, data nodes) are reordered with the recursive bisection heuristic RECBIS described in Section 2, as it has been proved to be very efficient for minimizing gaps between consecutive vertices in the adjacency lists [13]. In the same paper, the authors briefly mentioned a few ideas to improve RECBIS. Here, we propose some modifications in the direction that they suggest and incorporate two mechanisms to RECBIS to refine V_{\top} ordering π .

Partition swapping. The first mechanism is that of *partition swapping*. Let us recall that in RECBIS heuristic, a subset of nodes to reorder is partitioned into two sets V_1 and V_2 at each step. This bisection aims at minimizing the inner gaps of V_1 and V_2 . In the original paper, the method always assumes that nodes of V_1 precede those of V_2 in the final ordering π . It is however possible that swapping V_1 and V_2 leads to a better ordering because it can potentially imply smaller gaps between the partitions. We propose a linear-time heuristic to decide whether to apply the swap or not. The heuristic relies on the cost function defined below.

Definition 1 For $q \in V_{\perp}$, $cost_{ij}(q)$ be the gap-cost of q induced by the ordering $V_i V_j$ where $i, j \in \{1, 2\}, i \neq j$.

Note that the gap-cost is computed by summing over the \perp nodes (*i.e.*, the queries), as we aim to minimize the gaps between \top nodes, which are neighbors

³For bipartite data, locality is not relevant as a vertex is connected to vertices of the other group only.

of a same \perp node. At each step of the heuristic, we compute $\sum_{q \in V_\perp} cost_{12}(q)$ and $\sum_{q \in V_\perp} cost_{21}(q)$ and select the ordering between partitions that yields the lowest gap cost. To calculate the costs, we proceed as follows: given V_1 and V_2 , we define $\zeta_i(d)$ to be the position of d in the inner ordering of V_i and thus a number between 1 and $|V_i|$. Once the inner orderings of vertices inside V_1 and V_2 have been computed, we compute for each $q \in V_\perp$ and $i = 1, 2$ the quantities $max_i(q) = \max\{\zeta_i(d) : (q, d) \in E, d \in V_i\}$ and $min_i(q) = \min\{\zeta_i(d) : (q, d) \in E, d \in V_i\}$. These are respectively the largest and smallest ζ values of a neighbor of q in V_1 or V_2 , *i.e.*, the first and last positions a neighbor of q appears in V_i . Observe now that any difference between $cost_{12}(q)$ and $cost_{21}(q)$ must be due to the gap between q 's last appearance in the preceding partition and its first appearance in the following partition, that is to say the gap between $max_1(q)$ and $min_2(q)$ for V_1V_2 or $max_2(q)$ and $min_1(q)$ for V_2V_1 . We illustrate this idea in Figure 1. The formula for $cost_{ij}(q)$ is formally given by

$$cost_{ij}(q) = |V_i| + min_j(q) - max_i(q). \quad (2)$$




Figure 1: An example of the partition swapping heuristic for a \perp vertex q . If V_1 precedes V_2 , the gap associated with q is the distance between $max_1(q)$ and $min_2(q)$ *i.e.*, 3.

Partition flipping. To further improve the performance of the RECBIS heuristic, we also incorporate an additional mechanism that we call *partition flipping*. A flip in V_i is the act of reversing the order of the vertices in a partition V_i so that if the vertices are numbered from 1 to k then the vertex in position k' swaps position with the vertex in $k - k'$ for all $k' \in \llbracket 0, k - 1 \rrbracket$. Note that since flipping only reverses the order inside a partition, the gaps within a partition remain unchanged, but flipping allows to achieve smaller gaps between the partitions. For instance, considering the example of Figure 1, if we flip partition V_1 then the gap related to node q drops from 3 to 2, as the neighbor of q previously ordered first in V_1 is now ordered last in V_1 .

Adding both the swapping and flipping heuristics to RECBIS implies a linear-time overhead, as we simply need to iterate over the current neighbors of each $q \in V_\perp$ to calculate the $min_i(q)$ and $max_i(q)$ values. In practice, we observe that the overall time cost of these heuristics is negligible in comparison to the computation time of RECBIS while these heuristics bring an improvement to the compression rate that is typically of the order of 1%, as we will see in Section 4.2.

3.2 Bottom nodes ordering to improve referencing

We describe here our method to reorder \perp nodes (*i.e.*, query nodes according to our convention). Let us recall that the referencing mechanism consists in allowing a vertex to encode part of its neighborhood implicitly by representing it as a bit array which contains 1 for neighbors in common with one of its predecessors, called the reference. Consequently, the extent to which referencing is beneficial to compression largely depends on the ordering ϕ of the bottom nodes.

Efficient reordering approaches for \top nodes focus primarily on orderings that minimize gaps between adjacent vertices stored in the adjacency lists through the BIMLOGGAPA objective function, thus optimizing the delta encoding applied next. Here, we are rather looking for referencing-friendly orderings for \perp nodes *i.e.*, orderings which decrease storage costs through the efficient use of referencing. Note that orderings guaranteeing small gaps are likely to improve referencing too, however we argue that a V_{\perp} reordering method should be explicitly designed to favor efficient referencing. Ideally, we would like to express the reordering of V_{\perp} with an objective function designed in the same way that BIMLOGGAPA is to minimize gaps between V_{\top} . Unfortunately, the cost of referencing depends not only on the referencing scheme considered, but also on the scheme used to encode integers, and therefore on the compression software used following the reordering phase, such as WebGraph or Zuckerli (see [6, 29]). To explore the capabilities of the dual reordering scheme, we propose in what follows a software-agnostic optimization function for \perp nodes based on neighborhood similarity.

SIMREF: Improving referencing through similarity. We define here the SIMREF heuristic, which takes inspiration from the shingle heuristic [10]. Recall that the shingle is made to maximize the overlap between the neighborhoods of consecutive vertices. So, it should be efficient to order \perp nodes in line with favoring referencing. However, we observed in practice that it is not as effective as expected in this context. One reason for this shortcoming is that many nodes can share the same shingle, thus creating buckets of similar values, and vertices in a same bucket are ordered arbitrarily. As a result, it can happen that vertices with nearly identical neighborhoods are placed far apart in the same bucket, which can severely reduce the benefits of referencing.

The SIMREF heuristic orders the vertices within a bucket in a way that aims at maximizing the reference gain. The process works iteratively, by setting the position of one node in a bucket at each step. More precisely, assuming that a bucket contains at step i the ordered vertices b_1, \dots, b_{i-1} while c_i, c_{i+1}, \dots, c_k vertices are still unordered, the vertex from c_i, \dots, c_k selected to be in position b_i is the most similar to b_{i-1} . Similarity is measured with the Jaccard index, we remind here that $J(b, c) = |\mathcal{N}(b) \cap \mathcal{N}(c)| / |\mathcal{N}(b) \cup \mathcal{N}(c)|$, which should ensure to have a good overlap between the neighborhoods of two consecutive \perp nodes and thus favor an effective referencing. We describe the process more formally in Algorithm 1 for a bucket of size k corresponding to vertices with a same shingle.

Algorithm 1 Description of the SIMREF heuristic.

- 1: **Input:** randomly ordered table c_1, \dots, c_k of vertices with the same shingle
 - 2: **Output:** ordered table b_1, \dots, b_k
 - 3: **Initialize** $b_1 = c_1$
 - 4: **for** $i = 2$ **to** k **do**
 - 5: Compute $j_{max} = \operatorname{argmax}_{j \in [i, k]} (J(b_{i-1}, c_j))$
 - 6: Update $b_i \leftarrow c_{j_{max}}, c_{j_{max}} \leftarrow c_i$
-

As we will see in Section 4.3, the SIMREF heuristic brings improvements to the compression rate which vary depending on the dataset but can reach up to 6–7%. However, it is more computationally demanding than shingle is, as it adds a computation step which is quadratic in the size of the bucket⁴. Let us recall that the reordering time is not critical in general as this procedure is done offline and only once.

Nevertheless, it is possible to tune the trade-off between the computational cost of ordering nodes *vs.* the quality of the ordering for referencing purposes. Indeed, the heuristic can be adapted to compute the similarity of a subset of the candidate set $\mathcal{C} = \{c_i, \dots, c_j\}$, we briefly hint at two approaches for doing this. The first one is to sort the vertices in \mathcal{C} by degree, and select nodes which have the closest degree to the one of b_{i-1} , as their Jaccard similarity is more likely to be high. The second one is more elaborate, we generate a graph of nearest neighbors of the nodes in a bucket: we first compute for each vertex its k most similar neighbors where k is a parameter set by the user, then we create a graph G_s where the edge (u, v) exists iff u is one of v 's most similar neighbors. While considering b_{i-1} , its successor is selected among the available vertices in its neighborhood in G_s . Computing the k -nearest neighbors can be done with approximate algorithms, which exhibit good performance in nearly linear time [14].

4 Experiments

As described in Section 3, the dual ordering scheme that we propose consists in applying first the RECBIS procedure with the swapping and flipping improvements to reorder V_{\top} nodes by reducing the BIMLOGGAPA objective. Then we apply the SIMREF heuristic to reorder V_{\perp} nodes to improve the referencing. We describe in this Section the experiments carried out to evaluate the efficiency of this scheme.

⁴A reader could notice that the general problem may be seen as an instance of the well-known Traveling Salesman Problem (TSP) known to be NP-hard [18]: given a weighted graph, the TSP aims at finding a route of minimum cost that starts and ends in the same vertex and passes through all other vertices exactly once. Here the weight of edge (u, v) would correspond to $J(u, v)$, and SIMREF is a local search heuristic to approach a solution.

4.1 Data and protocol

We test our method on several massive bipartite graphs extracted from the KONECT⁵ network collection [20]. To evaluate the effect of the compression scheme in various contexts, we have selected several types of real-world datasets which can be represented by bipartite graphs. We first consider a *social network*, containing actors starring in movies extracted from the Internet Movie Database. Then we consider an *inverted index* which links together texts to the words that appear in them. Finally, we select a group of networks which represent *human online activity*: editing activity on a Wikipedia or a Wiktionary, user tagging songs on Delicious, user listening to songs on LastFM. When necessary, a dataset is pre-processed to eliminate multi-edges: for instance if a user has made several edits to a same Wikipedia page, it will be considered as one edge in the bipartite graph. The basic features of the datasets are described in Table 1.

Table 1: Number of nodes and edges of the bipartite graphs considered from [20].

Graph	V_{\top} nodes		V_{\perp} nodes		edges
imdb	303,617	movies	896,302	actors	3,782,463
Reuters	781,265	texts	283,911	words	60,569,726
lastfm-songs	992	users	1,084,620	songs	4,413,834
Delicious (user-tag)	833,081	users	4,512,099	tags	81,989,133
En-wiktionary-edits	66,140	users	5,826,113	pages	27,120,425
Fr-wikipedia-edits	757,621	users	8,870,762	pages	52,950,008
De-wikipedia-edits	1,025,084	users	5,910,432	pages	55,231,903

As a benchmark for comparison, we use the recursive bisection algorithm as proposed in [13] on bipartite data. Because RECBIS is not described in details for bipartite graphs in this paper, we consider two options: i) in the first one, it orders V_{\top} and V_{\perp} as if they were one set of nodes *i.e.*, the graph is processed in the same way as a unipartite graph would be, it is denoted RECBIS-u and it is the closest to the process described in [13]; ii) in the second one, the recursive bisection method is applied to V_{\top} and V_{\perp} separately, in the same way as our dual scheme works, it is denoted RECBIS-b. As we will see, both options roughly yield the same compression rates (on average, RECBIS-b very slightly outperforms RECBIS-u). In all cases, the resulting orderings serve as inputs to a standard compression method. The compression itself is achieved with Zuckerli [29] with its default settings, as it is the current state-of-the-art solution for graph compression. As is usual in the domain, the compression quality is measured in average number of bits per edge in the compressed graph.

4.2 RECBIS heuristics experiments

In the first set of experiments, we test the impact of the swapping and flipping heuristics from Section 3.1 on the recursive bisection method. The results with

⁵<http://konect.cc/>

RECBIS-b are summarized in Table 2. We can observe that in all tested instances there is a gain on the average number of bits per edge required to store the graph, but these gains remain rather marginal as they are typically of the order of 1% in the datasets that we have considered. However, as the time cost of these heuristics is only a small fraction of the overall time required to run RECBIS-b, we can recommend to systematically use the swapping and flipping heuristics when applying RECBIS.

Table 2: Average number of bits per edge in the compressed representation using the standard RECBIS-b algorithm derived from [13] and the version of the algorithm with the swapping and flipping heuristics (S&F).

Graph	RECBIS-b	RECBIS-b+S&F	gain (%)
imdb	10.30	10.15	1.46
Reuters	4.69	4.66	0.64
lastfm-songs	5.06	4.99	1.38
Delicious (user-tag)	6.86	6.82	0.58
En-wiktionary-edits	2.02	2.00	1.00
Fr-wikipedia-edits	6.54	6.48	0.92
De-wikipedia-edits	8.48	8.39	1.06

4.3 Dual reordering scheme experiments

We now present experimental results to evaluate the practical efficiency of the whole dual reordering scheme. The results are exhibited in Table 3. The Natural ordering denotes the compression obtained using the initial ordering of the nodes in the V_{\perp} as produced by the data collection method from [20] and RECBIS on those of V_{\top} and denotes a compression approach that does not make any attempt to optimize referencing. RECBIS-u corresponds to the baseline described in [13], where the vertices of V_{\perp} and V_{\top} are reordered together as if the graph were unipartite. The results with these two methods may be seen as standard approaches in the sense that they do not treat V_{\perp} and V_{\top} as different entities whereas the following two approaches do so. RECBIS-b denotes the approach which applies RECBIS separately on the sets of V_{\perp} and V_{\top} . Finally, our complete dual reordering scheme applying RECBIS on V_{\top} and SIMREF on V_{\perp} is denoted DUAL. Note that we want to separate the improvements due to the S&F heuristics from the improvement due to SIMREF, so in all uses of the RECBIS algorithm, we apply the proposed heuristic from Section 3.1, which implies that the RECBIS-b column is identical to the RECBIS-b+S&F heuristics column from Table 2.

Unsurprisingly, the worst compression rates are obtained with the Natural orderings, as there is no specific effort to optimize referencing. The performance of this method showcases the importance of reordering vertices to achieve efficient referencing. The results obtained with RECBIS-u and RECBIS-b are nearly identical. This observation stems from the fact that \top nodes (and respectively \perp nodes) are more similar to each other than they are to nodes of the

Table 3: Compression results using Zuckerli over different node orderings. The gain is the improvement from the DUAL reordering scheme to RECBIS-u (with the switching and flipping heuristics in all cases).

Graph	Natural	RECBIS-u [13]	RECBIS-b	DUAL	gain (%)
imdb	12.77	10.19	10.15	9.53	6.48
Reuters	4.71	4.67	4.66	4.66	0.21
lastfm-songs	6.00	4.98	4.99	4.66	6.43
Delicious (user-tag)	7.53	6.82	6.82	6.69	1.91
En-wiktionary-edits	4.11	2.00	2.00	1.91	4.50
Fr-wikipedia-edits	8.89	6.46	6.48	6.38	1.24
De-wikipedia-edits	10.24	8.39	8.39	8.28	1.31

other group, consequently RECBIS-u tends to not intermix the two groups of nodes and thus orders nodes as RECBIS-b does. Most importantly, we observe that the compression rates obtained with the dual reordering scheme (DUAL column) outperforms the other methods. The gain depends on the dataset under consideration being lower than 2% for some datasets (**Reuters**, **Delicious**, **Fr/De-wikipedia-edits**) but reaching up to 6 – 7% for others (**imdb**, **lastfm**), which is substantial in the domain of lossless graph compression.

A few additional remarks regarding these results guide us to look for further improvements. First, the two ordering steps of the dual scheme can be iterated to compress further the graph until the orderings do not significantly change. However, we observed that the improvements brought by iterating the process are marginal on the datasets under study (typically less than 1%, not reported here). Second, it is possible to switch the roles of \top and \perp nodes in the networks that we considered, as we are representing them as bipartite graphs without regard for what the nodes represent. The results that we have shown here correspond to the choice yielding the best compression rates for each dataset. We can see in Table 1 that in all cases (except for **Reuters**, which always yields poor gains), $|V_{\top}| < |V_{\perp}|$. It seems to indicate that improving referencing yields better results when there is a larger choice of reference nodes to pick from.

From our perspective, the most important thing to draw from these experiments is that the DUAL ordering scheme seems promising as it consistently outperforms other ordering methods, and that this mainly stems from the improvement made on the ordering ϕ which targets referencing. We can thus think of further developments to improve the referencing gains in bipartite graphs, as discussed in the next section.

5 Discussion on further developments

Combining ordering heuristics. We have implemented the dual reordering scheme on bipartite graphs using a combination of RECBIS for ordering \top nodes and SIMREF for ordering \perp nodes with encouraging results. It is likely that other heuristic combinations can lead to better compression rates. For

instance, SIMREF is based on the Jaccard similarity, but other vertex similarity measures can be implemented, such as the Adamic-Adar index [1] or the Resource Allocation index [30].

The efficiency of the combinations is certainly data-dependent. Understanding precisely why a dataset benefits more or less from a compression scheme originates from the particularities of the graph. This point deserves deeper examination, and we think that the tools used in Zuckerli [29] to investigate which parts of the compressed graph require the most bits is useful to pinpoint steps in the compression process that could be improved.

Referencing scheme. Another lead is to investigate the referencing scheme itself and possibly define a similarity metric based on this scheme. We suggest here two ideas which can be interesting ways to improve the referencing scheme.

First, we can apply a *post-processing reordering* on the ϕ ordering for \perp nodes. Supposing that we represent the fact that u references v by a directed arc from u to v , then the references among \perp nodes form a forest: each vertex can only have one reference – its parent – and can be referenced by several nodes – its children. In addition, when vertex u references another vertex v , the gap between u and v must be stored. Given such a tree, we can create a new ordering ϕ' where nodes of V_\perp in the same tree of the forest, are given consecutive ids. This reordering will prohibit vertices in different trees from alternating with each other in the final ordering and thus decrease the referencing gaps between them, thereby improving the final compression rate.

Another possibility is to allow *forward referencing*. With the usual referencing techniques discussed earlier, a node can only select its reference backwards, in the sense that we look for a reference in a window of nodes located before in the ϕ ordering. By allowing a node to select its reference forward, we allow the referencing scheme to be more flexible, which can potentially improve the compression rates. To illustrate this idea, let us consider a toy example of two nodes $u, v \in V_\perp$, u has the following neighbors: $2, 4, \dots, 10$, while v has $1, 2, 3, \dots, 10$, in addition u precedes v in ϕ . With a standard referencing, we can only have $u \rightarrow v$ which leads to $1, 3, \dots, 9$ being residual neighbors that are encoded separately with delta encoding. As the neighbors of v can be encoded more efficiently because they are consecutive, this reference would not be picked. With the forward scheme, the reference $v \rightarrow u$ is also considered, and it would be efficient as u could have its neighbors encoded with a simple bit array. Note however, that this modification comes at a cost: references no longer form a tree but a directed graph, and we must ensure that the directed graph of references is a directed acyclic graph, *i.e.*, there is no cycle of references.

Preliminary results – based on real data but only evaluating the expected referencing gain – indicate that these modifications could improve the compression rate up to a few percents. To go further in this direction, it is necessary to select a specific compression method and to dive deep into its code to adjust the methods to its specificities, which we leave to future investigations.

6 Conclusion

In this work, we have examined the problem of lossless compression of bipartite graphs and proposed a dual reordering scheme of the vertices. The central idea is to reorder the vertices of each partition with a different perspective in mind: either to optimize delta encoding or to maximize the effect of referencing, two techniques which are essential for standard compression methods. We have shown empirically that this approach outperforms the classic single ordering methods, however the range of the improvement varies significantly depending on the dataset under study. These encouraging results guided us to propose several leads for further improvements with the idea that reordering can be suited to specific datasets and specific referencing schemes.

Acknowledgements

We thank Fabrice Lécuyer and Matthieu Latapy for their proofreading of the manuscript. This work is funded by the ANR (French National Agency of Research) partly by the Limass project (under grant ANR-19-CE23-0010).

References

- [1] Adamic, L.A., Adar, E.: Friends and neighbors on the web. *Social networks* **25**(3), 211–230 (2003)
- [2] Apostolico, A., Drovandi, G.: Graph compression by bfs. *Algorithms* **2**(3), 1031–1044 (2009)
- [3] Besta, M., Hoefler, T.: Survey and taxonomy of lossless graph compression and space-efficient graph representations. *arXiv preprint arXiv:1806.01799* (2018)
- [4] Block, P.: Reciprocity, transitivity, and the mysterious three-cycle. *Social Networks* **40**, 163–173 (2015)
- [5] Boldi, P., Rosa, M., Santini, M., Vigna, S.: Layered label propagation: A multiresolution coordinate-free ordering for compressing social networks. In: *Proceedings of the 20th international conference on World Wide Web*. pp. 587–596 (2011)
- [6] Boldi, P., Vigna, S.: The webgraph framework i: compression techniques. In: *Proceedings of the 13th international conference on World Wide Web*. pp. 595–602 (2004)
- [7] Brisaboa, N.R., Ladra, S., Navarro, G.: k 2-trees for compact web graph representation. In: *International symposium on string processing and information retrieval*. pp. 18–30. Springer (2009)

- [8] Broder, A.Z., Charikar, M., Frieze, A.M., Mitzenmacher, M.: Min-wise independent permutations. *Journal of Computer and System Sciences* **60**(3), 630–659 (2000)
- [9] Buehrer, G., Chellapilla, K.: A scalable pattern mining approach to web graph compression with communities. In: *Proceedings of the 2008 international conference on web search and data mining*. pp. 95–106 (2008)
- [10] Chierichetti, F., Kumar, R., Lattanzi, S., Mitzenmacher, M., Panconesi, A., Raghavan, P.: On compressing social networks. In: *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. pp. 219–228 (2009)
- [11] Claude, F., Navarro, G.: Fast and compact web graph representations. *ACM Transactions on the Web (TWEB)* **4**(4), 1–31 (2010)
- [12] Delpratt, O., Rahman, N., Raman, R.: Engineering the louds succinct tree representation. In: *International Workshop on Experimental and Efficient Algorithms*. pp. 134–145. Springer (2006)
- [13] Dhulipala, L., Kabiljo, I., Karrer, B., Ottaviano, G., Pupyrev, S., Shalita, A.: Compressing graphs and indexes with recursive graph bisection. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 1535–1544 (2016)
- [14] Dong, W., Charikar, M., Li, K.: Efficient k-nearest neighbor graph construction for generic similarity measures. In: *Proceedings of the 20th international conference on World wide web*. pp. 577–586 (2011)
- [15] Elias, P.: Universal codeword sets and representations of the integers. *IEEE transactions on information theory* **21**(2), 194–203 (1975)
- [16] Fiduccia, C.M., Mattheyses, R.M.: A linear-time heuristic for improving network partitions. In: *19th design automation conference*. pp. 175–181. IEEE (1982)
- [17] Grabowski, S., Bieniecki, W.: Tight and simple web graph compression for forward and reverse neighbor queries. *Discrete Applied Mathematics* **163**, 298–306 (2014)
- [18] Jünger, M., Reinelt, G., Rinaldi, G.: The traveling salesman problem. *Handbooks in operations research and management science* **7**, 225–330 (1995)
- [19] Kernighan, B.W., Lin, S.: An efficient heuristic procedure for partitioning graphs. *The Bell system technical journal* **49**(2), 291–307 (1970)
- [20] Kunegis, J.: KONECT – The Koblenz Network Collection. In: *Proc. Int. Conf. on World Wide Web Companion*. pp. 1343–1350 (2013), <http://dl.acm.org/citation.cfm?id=2488173>

- [21] Larsson, N.J., , Moffat, A.: Off-line dictionary-based compression. *Proceedings of the IEEE* **88**(11), 1722–1732 (2000)
- [22] Liakos, P., Papakonstantinou, K., M. Sioutis, M.: Pushing the envelope in graph compression. In: *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. pp. 1549–1558 (2014)
- [23] Moffat, A.: Huffman coding. *ACM Computing Surveys (CSUR)* **52**(4), 1–35 (2019)
- [24] Moffat, A., Neal, R.M., Witten, I.H.: Arithmetic coding revisited. *ACM Transactions on Information Systems (TOIS)* **16**(3), 256–294 (1998)
- [25] Pibiri, G.E., Venturini, R.R.: Techniques for inverted index compression. *ACM Computing Surveys (CSUR)* **53**(6), 1–36 (2020)
- [26] Pibiri, G.E., Petri, M., Moffat, A.: Fast dictionary-based compression for inverted indexes. In: *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. pp. 6–14 (2019)
- [27] Rice, R., Plaunt, J.: Adaptive variable-length coding for efficient compression of spacecraft television data. *IEEE Transactions on Communication Technology* **19**(6), 889–897 (1971)
- [28] Rossi, R.A., Zhou, R.: Graphzip: a clique-based sparse graph compression method. *Journal of Big Data* **5**(1), 1–14 (2018)
- [29] Versari, L., Comsa, I.M., Conte, A., Grossi, R.: Zuckerli: A new compressed representation for graphs. *IEEE Access* **8**, 219233–219243 (2020)
- [30] Zhou, T., Lü, L., Zhang, Y.C.: Predicting missing links via local information. *The European Physical Journal B* **71**(4), 623–630 (2009)