

# Prefix-Free Parsing for Building Large Tunnelled Wheeler Graphs

Adrián Goga<sup>1</sup> ✉

Department of Computer Science, Faculty of Mathematics, Physics and Informatics,  
Comenius University, Bratislava, Slovakia

Andrej Baláž ✉

Department of Applied Informatics, Faculty of Mathematics, Physics and Informatics,  
Comenius University, Bratislava, Slovakia

---

## Abstract

We propose a new technique for creating a space-efficient index for large repetitive text collections, such as pangenomic databases containing sequences of many individuals from the same species. We combine two recent techniques from this area: Wheeler graphs (Gagie et al., 2017) and prefix-free parsing (PFP, Boucher et al., 2019).

Wheeler graphs are a general framework encompassing several indexes based on the Burrows-Wheeler transform (BWT), such as the FM-index. Wheeler graphs admit a succinct representation which can be further compacted by employing the idea of tunnelling, which exploits redundancies in the form of parallel, equally-labelled paths called blocks that can be merged into a single path. The problem of finding the optimal set of blocks for tunnelling, i.e. the one that minimizes the size of the resulting Wheeler graph, is known to be NP-complete and remains the most computationally challenging part of the tunnelling process.

To find an adequate set of blocks in less time, we propose a new method based on the prefix-free parsing (PFP). The idea of PFP is to divide the input text into phrases of roughly equal sizes that overlap by a fixed number of characters. The phrases are then sorted lexicographically. The original text is represented by a sequence of phrase ranks (the parse) and a list of all used phrases (the dictionary). In repetitive texts, the PFP representation of the text is generally much shorter than the original since individual phrases are used many times in the parse, thus reducing the size of the dictionary.

To speed up the block selection for tunnelling, we apply the PFP to obtain the parse and the dictionary of the original text, tunnel the Wheeler graph of the parse using existing heuristics and subsequently use this tunnelled parse to construct a compact Wheeler graph of the original text. Compared with constructing a Wheeler graph from the original text without PFP, our method is much faster and uses less memory on collections of pangenomic sequences. Therefore, our method enables the use of Wheeler graphs as a pangenomic reference for real-world pangenomic datasets.

**2012 ACM Subject Classification** Theory of computation → Theory and algorithms for application domains

**Keywords and phrases** Wheeler graphs, BWT tunnelling, prefix-free parsing, pangenomic graphs

**Digital Object Identifier** 10.4230/LIPIcs.WABI.2022.18

**Related Version** *arXiv Version*: <https://arxiv.org/abs/2206.15097>

**Supplementary Material** *Software (Source Code)*: [https://github.com/fmfi-compbio/pfp\\_wg](https://github.com/fmfi-compbio/pfp_wg)  
archived at `swh:1:dir:580f9ea6d61e879c69ebb8a8337902e77a3f8a2b`

**Funding** *Adrián Goga*: VEGA grant 1/0463/20; EU Horizon 2020 grant No. 956229 (ALPACA); Comenius University grant for doctoral students No. 422

*Andrej Baláž*: VEGA grant 1/0538/22; EU Horizon 2020 grant No. 956229 (ALPACA)

---

<sup>1</sup> Corresponding author



**Acknowledgements** We want to thank Travis Gagie for the conception of the idea during his data structures course and helpful remarks throughout the realisation of this project. Our thanks also go to Broňa Brejová for useful advice during the writing process and Uwe Baier for kindly responding to our questions via email. Finally, we thank Lucas Pansani Ramos for the aid he provided in the early days of the project.

## 1 Introduction

The discovery of Burrows-Wheeler transformation (BWT) [6], a text permutation initially intended for data compression, and the following breakthrough that enhanced it with indexing properties started a new paradigm of text indexing. In this paradigm, the input text can be stored in space not far larger than its empirical entropy while allowing for efficient pattern matching queries. The essence of the BWT, the idea of suffix sorting, has been since enjoyed by a large variety of popular genomic indexing tools. Sequencing read alignment tools, such as BWA [16], Bowtie [14, 15] or CHIC [23] are directly based upon the BWT, while others such as VG [10] use some of its aspects.

As the BWT is an essential tool in processing large textual datasets, it motivates us to further minimize its space requirements. Since its conception, a popular choice has been to use move-to-front (MTF) transformation and subsequently the run-length encoding (RLE), followed by an entropy encoding, e.g. Huffman [6]. Quite recently, in 2018, Baier [2] described another novel source of redundancy in BWTs, yet undetected by the RLE, and proposed a technique called *tunnelling* purposed to deal with this particular redundancy. Tunnelled BWTs are a case of Wheeler graphs [9], a general framework of the text indices based on the idea of suffix sorting.

Tunnelling has been experimentally proven to reduce the size of the resulting BWTs, especially in repetitive texts. The size of the input files is a significant bottleneck of this approach for two reasons. The first one is the fast BWT construction algorithm which requires linear, although non-negligible, memory overhead over the input file itself. The second reason is that tunnelling requires random access to the BWT, rendering it impractical when working in external memory.

The first of these issues could be alleviated by employing a suitable BWT construction algorithm, e.g. one that creates the BWT from a compressed representation of the input, such as the Lempel-Ziv parse by Policriti and Prezza [20] or the Prefix-Free Parsing (PFP) method by Boucher et al. [5]. PFP is especially interesting due to its simplicity and effectiveness, accomplishing to represent large repetitive texts in orders of magnitude smaller space [5]. The second issue is harder to tackle, requiring the attention to either adapt the particular tunnelling algorithm to find the tunnelled BWT from the compressed representation or, when possible, tunnelling the BWT of the compressed representation and then using the information to output a tunnelled BWT of the original input.

The objective of our work is to resolve both problems simultaneously. We have adapted the PFP approach for building BWTs from large repetitive datasets in such a way that it produces tunnelled BWTs. While the resulting BWTs are larger than the original tunnelled BWTs, our approach scales to much larger dataset sizes and could be improved by further post-processing. Our method is not limited to producing tunnelled BWTs, but virtually any Wheeler graphs and hence provides a good starting point for building large pangenomic indexes.

## 1.1 Preliminaries

In this work, we will use the following notation. Let  $\Sigma$  be the working alphabet, i.e. a finite set of symbols totally ordered by  $\preceq$ . We denote the size of  $\Sigma$  as  $|\Sigma| = \sigma$ . A *string*  $T = T[1] \dots T[n]$  is a sequence of symbols from  $\Sigma$ . We use the terms *string* and *text* interchangeably, similarly with *symbol* and *character*. A string made of a single character is called *unary*. Furthermore, strings that end with  $\$$  are called *null-terminated*, where  $\$$  is the lexicographically smallest symbol in  $\Sigma$  (i.e.  $\$ \preceq c$  for any  $c \in \Sigma$ ) that may only appear at the right end. Moreover, we extend the use of  $\preceq$  to label the lexicographic order of strings and write  $\alpha \preceq \beta$  if and only if  $\alpha$  is lexicographically smaller or equal than  $\beta$ , where  $\alpha, \beta \in \Sigma^*$ . We will write  $\alpha < \beta$  if  $\alpha \preceq \beta$  and  $\alpha \neq \beta$ .


The length of the string  $T$  will be denoted by  $|T| = n$ . The string  $\varepsilon$  is the unique string of length 0. A character of  $T$  at position  $i$  is denoted by  $T[i]$ . A substring  $T[i \dots j]$  of  $S$  is a subsequence  $T[i]T[i+1] \dots T[j]$ . We assume  $T[i \dots j] = \varepsilon$  if  $i > j$ . The substring  $T[1 \dots j]$  is called the  $j$ -th *prefix* of  $T$ , while the substring  $T[j \dots n]$  is called the  $j$ -th *suffix*.

For integers  $i, j$  we define the set  $\{k \mid k \in \mathbb{N}; i \leq k \leq j\}$  to be the *interval* between  $i$  and  $j$  and denote it as  $[i, j]$ .

## 2 Background

### 2.1 Burrows-Wheeler transform

The Burrows-Wheeler transform (BWT) [6] is an incredibly influential transformation of text, formed by the concatenation of the letters in the last column of the Burrows-Wheeler matrix, which is created by lexicographically sorting all possible rotations of the original text. This construction is visualized in Figure 1.



**Figure 1** BWT construction. All rotations of the string  $ACGTCGTT\$$  are sorted lexicographically in the Burrows-Wheeler matrix. The last column of Burrows-Wheeler matrix represents the BWT.

The importance of the BWT in string processing is derived from its compressibility, reversibility and usefulness in pattern matching. At first, the improvement in the compressibility of BWT in comparison with the original string is not apparent. However, it stems from empirical observation that similar contexts tend to be preceded by the same characters. The construction of BWT sorts these contexts and therefore tends to place identical characters in runs, which improves the compression.

The reversibility of the BWT is guaranteed by the so-called LF mapping, which uses the observation that the ranks of letters remain the same between the BWT and the first column of the Wheeler matrix F (i.e. the first letter in BWT corresponds to the first letter in F). Since column F can be recovered from BWT by simply sorting the letters of BWT,

the original text can be reconstructed in reverse order by following the LF mapping from the first row of the implicit Wheeler matrix. Furthermore, the LF mapping, together with the suffix array, can be used for efficient pattern matching in the data structure known as FM-index [7].

Although illustrative, the aforementioned construction of BWT is time-inefficient and seldom used in practice. Several efficient methods for construction in linear time exist [13, 5]. This work utilizes prefix-free parsing (PFP), a technique first used for building BWTs in [5]. In comparison with previous methods, the PFP approach allowed the construction of BWT in sublinear space for repetitive data and therefore unlocked a possibility to produce BWTs of big datasets, of which sizes vastly exceed the memory limits of current computers.

## 2.2 Prefix-free parsing

In prefix-free parsing, the original text  $T$  is divided into phrases of variable length, which are stored in the dictionary  $\mathcal{D}$ . Furthermore, the parse  $\mathcal{P}$  is created as a list of lexicographic ranks of phrases, in order in which the phrases appear in the original text. The parse  $\mathcal{P}$  and the dictionary  $\mathcal{D}$  allow us to reconstruct the original text  $T$ .

To perform the division of  $T$  into phrases, we define a set of trigger words  $E$ , where each trigger word is of length  $w$ . Subsequently, we use a sliding window of size  $w$  through the text  $T$  and each time the sliding window matches a trigger word in  $E$ , we terminate the current phrase, add it to the dictionary, and initiate a new phrase. In the end, the dictionary is sorted lexicographically, and the parse is relabeled accordingly. It is noteworthy that the phrases in the dictionary begin and end with a trigger word, and consequently, no phrase is a prefix of another phrase in the dictionary.

This PFP construction is valuable for several reasons. Firstly, only a single pass through the original text is needed to obtain all phrases, and therefore the text can be sequentially read from the disk allowing the processing of large texts.

Secondly, suppose the original text is repetitive. In that case, the parse and the dictionary tend to occupy a much smaller space than the raw representation, and therefore the PFP can be loaded into memory.

Finally, we can use a rolling hash  $h : \text{word} \rightarrow \{0 \dots p - 1\}$  to identify the trigger words. If the hash is equal to zero, we consider it a member of the set of trigger words  $E$ . The parameter  $p$  then allows us to adjust the lengths of phrases since the expected length of a phrase in a uniformly-random text is  $p$ .

An example of a prefix-free parse construction is shown in Example 1.

► **Example 1.** Let us have a text  $T = \#ABDACDABDACDA\$$ . For the purpose of this example, let the set of trigger words be  $E = \{A\}$ . Then we have  $D = \#A, ABDA, ACDA, A\$$  and  $\mathcal{P} = [0, 1, 2, 1, 2, 3]$ . Then we lexicographically sort the dictionary phrases to get  $\#A, A$, ABDA, ACDA$  and the remapped parse  $\mathcal{P} = [0, 2, 3, 2, 3, 1]$ .


## 2.3 Wheeler graphs

Wheeler graphs (WGs), introduced by Gagie et al. [9], are a class of labelled graphs that generalize the notion of BWT and several of its variations into a unified framework. A formal definition of Wheeler graphs is presented in Definition 2.

► **Definition 2 (Wheeler graph).** Let  $G = (V, E)$  be a directed graph labeled by  $\lambda : E \rightarrow \Sigma$ . Then  $G$  is a Wheeler graph if and only if there is a total order  $\leq$  of its vertices, also called the Wheeler order, such that for each pair of edges  $e_1 = (u_1, v_1), e_2 = (u_2, v_2)$  the following conditions hold:

1. The vertices with zero in-degree precede those with non-zero in-degree,
2. If  $\lambda(e_1) \prec \lambda(e_2)$ , then  $v_1 < v_2$ ,
3. If  $\lambda(e_1) = \lambda(e_2)$  and  $u_1 < u_2$ , then  $v_1 \leq v_2$ .

To illustrate this definition, we can visualize the Wheeler graphs by duplicating the vertices and displaying them in two columns in their Wheeler order, as shown in Figure 2. The edge  $(u, v)$  will be drawn from vertex  $u$  in the right column to vertex  $v$  in the left column. Then, the first condition guarantees that the vertices without an incoming edge will form a single interval in the ordering, and that interval will be placed at the top. The second condition guarantees that the vertices with incoming edges labelled by a character  $c$  will form an interval, and those intervals appear from top to bottom according to the lexicographic order of the character  $c$ . The third condition guarantees that no two edges with the same label  $c$  will cross.




■ **Figure 2** The Wheeler graph visualization. The vertices are doubled and shown in two columns to highlight the conditions of Definition 2. Visualization also reveals the connection between WGs and LF mapping in BWTs.

These conditions give rise to an essential property of WGs, called *path coherence*, which is a generalization of the property from BWTs that enables the LF mapping. Therefore, WGs can be used for efficient pattern matching and to store textual data in a lossless manner. Furthermore, WGs allow for a succinct representation by storing the following data.

- An array  $C$  of length  $\sigma$  s.t.  $C[c]$  is the number of edges labeled by characters lexicographically smaller than  $c$ ,
- The string  $L = L_1 \dots L_n$ , where  $L_i$  is the concatenation of the labels of the outgoing edges of vertex  $v_i$ ,
- The bit vector  $O = 10^{d_{out}(v_1)-1} \dots 10^{d_{out}(v_n)-1}$ ,
- The bit vector  $I = 10^{d_{in}(v_1)-1} \dots 10^{d_{in}(v_n)-1}$ ,

where  $d_{in}(v)$  ( $d_{out}(v)$ ) is the in-degree (out-degree) of the vertex  $v$ .

In comparison with the plain list-of-edges representation, which uses  $\Theta(|E| \log \sigma + |E| \log |V|)$  bits of space, the succinct representation uses only  $\Theta(|E| \log \sigma + 2|E| + \sigma \log |E|)$  and together with efficient rank and select queries on the bit vectors  $O$  and  $I$  supports efficient graph traversing.



■ **Figure 3** Tunnelling of the WG. Both paths  $8 - 5 - 3$  (green) and  $7 - 4 - 2$  (blue) spell **GC** and are merged into a single path in the reduced WG.

Another compelling property of WGs is their reducibility, where a WG representing a particular set of strings can be replaced by a smaller WG representing the same set of strings. One way to accomplish this is by treating the given WG as a Wheeler nondeterministic finite automaton (WNFA), which we subsequently convert to its deterministic variant (WDFA). It has been shown that such conversions for automata that correspond to WGs result only in a linear number of additional states [1]. For a WDFA, an equivalent WDFA with the minimum number of states can then be found by a linear-time algorithm by Alanko et al. [21].

Another particular process of reducing a WG is called *tunnelling* [2] and is illustrated in Figure 3. During the tunnelling, the paths  $p_1 = (v_{1,1}, \dots, v_{1,\ell}), \dots, p_k = (v_{k,1}, \dots, v_{k,\ell})$  of which the vertex sets  $V_j = \{v_{i,j} \mid 1 \leq i \leq k\}$  are adjacent in the Wheeler order for each  $1 \leq j \leq \ell$ , and which spell the same substring are identified and a tunnel is formed by merging the vertices from  $V_j$  into a single vertex  $v_j$ . In general, the task of finding the smallest WG via tunneling was proven to be NP-hard, and therefore multiple heuristic approaches to solve this problem were proposed [4].

### 3 Methods

We propose to employ PFP as a preprocessing step in building Wheeler graph indexes from large datasets through tunnelling BWTs. We are inspired by the original work by Boucher et al. [5] of building large BWTs using small working memory.

After obtaining the parse  $\mathcal{P}$  of a text  $T$ , we build the  $BWT(\mathcal{P})$  and subsequently find its tunnelling using one of the approaches proposed by Baier and Dede [4]. Since the parse  $\mathcal{P}$  is a sequence of lexicographic ranks of dictionary phrases, each character of the parse  $\mathcal{P}$  corresponds to multiple characters of  $T$  (controlled to a certain degree by the PFP parameter  $p$ ), and hence serves as a significantly shortened representation of  $T$ . Therefore, tunnelling  $BWT(\mathcal{P})$  is practical for much larger file sizes.

Under the additional assumption that the input files are repetitive, we can also deduce that at least a certain degree of repetitiveness will also be present in  $BWT(\mathcal{P})$ , and hence  $BWT(\mathcal{P})$  can be significantly shortened by tunnelling. Once we obtain a tunnelled  $BWT(\mathcal{P})$  in the form of a Wheeler graph  $G_{\mathcal{P}}$  of the parse  $\mathcal{P}$ , our goal is to efficiently transform it to the Wheeler graph  $G_T$  of the text  $T$  while retaining the tunnelled paths in an appropriate form.

### 3.1 Theory

We begin by introducing a vital claim that was already proven by Boucher et al. [5].

► **Definition 3** (Suffix set [5]). *Let  $T$  be a text of length  $n$ . We call the set  $Suf(T) = \{T[i \dots n] \mid i \in [1, n]\}$  a suffix set of  $T$ .*

► **Lemma 4.** *Let  $T$  be a text,  $\mathcal{P}$  be the parse of the PFP( $T$ ) and  $x, y \in Suf(T)$  such that  $x \prec y$ . Furthermore, let  $\alpha_x, \alpha_y$  be prefixes of  $x, y$  up to the nearest phrase boundaries and  $\beta_x, \beta_y$  be the rest of  $x, y$ . The  $\beta_x, \beta_y$  correspond to suffixes  $\gamma_x, \gamma_y$  in the parse  $\mathcal{P}$ . Then either it holds that  $\alpha_x \prec \alpha_y$  or  $\alpha_x = \alpha_y$  and  $\gamma_x \prec \gamma_y$ .*

**Proof.** The claim follows from the lemmas 1-7 proven by Boucher et al. [5]. ◀

The lemma 4 allows the  $BWT(T)$  to be constructed from the parse  $\mathcal{P}$  and the dictionary  $\mathcal{D}$  by iterating over the suffixes of the phrases from  $\mathcal{D} = \{D_1, \dots, D_k\}$  sorted in the lexicographic order and outputting the characters preceding their occurrences in  $T$ . As the vertices of Wheeler graphs are totally ordered according to the suffixes of  $T$  they represent, we immediately obtain a straightforward generalization of this result for Wheeler graphs in the sense that a Wheeler graph of  $\mathcal{P}$  implies a Wheeler graph of  $T$ . The exact meaning of this implication is detailed in Def. 5.

► **Definition 5** (Expanded graph). *Let  $T$  be a string and  $G_{\mathcal{P}} = (V_{\mathcal{P}} = [1, m], E_{\mathcal{P}}, \lambda_{\mathcal{P}})$  be a Wheeler graph of  $\mathcal{P}$  from the PFP( $T$ ), the vertices of which are labeled according to a Wheeler order. Then we define the graph  $G' = (V', E', \lambda')$  to be the Expanded graph of  $T$  constructed from  $G_{\mathcal{P}}$ , where*

$$\begin{aligned} V_{int} &= \{v_{e,2}, \dots, v_{e,\ell-w} \mid e \in E_{\mathcal{P}}, \ell = |D_{\lambda_{\mathcal{P}}(e)}|\}, \\ V' &= V_{\mathcal{P}} \cup V_{int}, \\ E' &= \{(v_{e,i}, v_{e,i-1}) \mid e \in E_{\mathcal{P}}, i \in [3, \ell - w], \ell = |D_{\lambda_{\mathcal{P}}(e)}|\} \cup \\ &\quad \{(t, v_{e,\ell-w}) \mid e = (t, u) \in E_{\mathcal{P}}, \ell = |D_{\lambda_{\mathcal{P}}(e)}|\} \cup \\ &\quad \{(v_{e,2}, u) \mid e = (t, u) \in E_{\mathcal{P}}\} \\ \lambda'(e') &= \begin{cases} D_{\lambda_{\mathcal{P}}(e)}[i-1] & \text{if } e' = (v_{e,i}, v_{e,i-1}), e \in E_{\mathcal{P}}, \ell = |D_{\lambda_{\mathcal{P}}(e)}|, i \in [3, \ell - w], \\ D_{\lambda_{\mathcal{P}}(e)}[\ell - w] & \text{if } e' = (t, v_{e,\ell-w}), e \in E_{\mathcal{P}}, \ell = |D_{\lambda_{\mathcal{P}}(e)}|, \\ D_{\lambda_{\mathcal{P}}(e)}[1] & \text{if } e' = (v_{e,2}, u), e = (v, u) \in E_{\mathcal{P}}. \end{cases} \end{aligned}$$

The expanded graph  $G_T = (V, E, \lambda)$  is therefore obtained by expanding each edge  $e = (t, u) \in E_{\mathcal{P}}$  labeled  $\lambda_{\mathcal{P}}(e)$  into a path with newly added internal vertices  $v_{e,2}, \dots, v_{e,\ell-w}$  so that the concatenation of the labels  $\lambda((v_{e,2}, u))\lambda((v_{e,3}, v_{e,2})) \dots \lambda((v_{e,\ell-w}, v_{e,\ell-w-1}))\lambda((t, v_{e,\ell-w})) = D_{\lambda_{\mathcal{P}}(e)}[1 \dots \ell - w]$ , where  $\ell = |D_{\lambda_{\mathcal{P}}(e)}|$ . In other words, if we take an edge  $e = (t, u) \in E_{\mathcal{P}}$  and an internal vertex



$v$  from the expansion of  $e$  in  $G_T$ , then starting from  $v$  and outputting the edge labels while traversing the reversed edges up to  $t$  we get a suffix of  $D_{\lambda(e)}[1 \dots \ell - w]$ . With the help of Def. 6, we extend this notion to non-internal vertices in Def. 7, which we then use in Thm. 8.

► **Definition 6** (Incoming label). *Let  $G = (V, E, \lambda)$  be a Wheeler graph. Then for any  $v \in V$  we will call the value  $i(v) = \lambda((u, v))$  for such  $u \in V$  that  $(u, v) \in E$  an incoming label of the vertex  $v$ .*

► **Definition 7** (Phrase suffix). *Let  $G_{\mathcal{P}} = (V_{\mathcal{P}}, E_{\mathcal{P}}, \lambda_{\mathcal{P}})$  be a Wheeler graph over a PFP of the string  $T$ . Let  $G = (V, E, \lambda)$  be the expanded graph of  $G_{\mathcal{P}}$ . Then for any  $v \in V$  we defined  $f(v)$  as*

$$f(v) = \begin{cases} D_{i(w)}[i \dots \ell - w], & \text{if } v = v_{e,i}, e = (u, w) \in E_{\mathcal{P}}, i \in [2, \ell - w], \ell = |D_{\lambda_{\mathcal{P}}(e)}| \\ D_{i(v)}[1 \dots \ell - w], & \text{otherwise} \end{cases}$$

Note that in a Wheeler graph, all the incoming edges to a vertex  $v$  have to be labelled the same, hence the correctness of Def. 6. Def. 7 deals with two cases – the first one being the internal vertex of the expanded Wheeler graph and a vertex of the  $G_{\mathcal{P}}$ . Now we are ready to state our core theorem.

► **Theorem 8.** *Let  $G_{\mathcal{P}} = (V_{\mathcal{P}}, E_{\mathcal{P}})$  be a Wheeler graph over a PFP of the string  $T$ . Let  $G = (V, E)$  be the expanded graph of  $G_{\mathcal{P}}$ . Then  $G$  is also a Wheeler graph.*

**Proof.** We prove the claim by showing that there is an ordering of vertices from  $V$  such that the conditions from Def. 2 are fulfilled. We define the order as follows. Let  $v_1, v_2 \in V$  and  $f(v_i)$  a phrase suffix according to Def. 7 for  $i = 1, 2$ . Then  $v_1 \prec v_2$  if and only if  $f(v_1) \prec f(v_2)$  or  $f(v_1) = f(v_2)$  and  $u_1 < u_2$ .

We have no vertices of zero in-degree in  $G$ , so property 1 is satisfied for any vertex order. We need to show that the vertices in  $V$  are sorted according to the suffixes of  $T$  they represent. However, this follows from the lemma 4 and the definition of our order. ◀

We note that finding an ordering of the nodes of Wheeler graphs is *NP*-complete in general [11]; however, it can indeed be accomplished in a polynomial time in our restricted case, in which we already have the relative order of the vertices with in-degree and out-degree larger than one<sup>2</sup>.

To build the expanded Wheeler graph, we follow an approach similar to that of Boucher et al. [5] for building large BWTs from  $\mathcal{P}$  and  $\mathcal{D}$ . First, we construct a suffix array  $SA_{\mathcal{D}}$  of the concatenation of all phrases from the dictionary  $\mathcal{D}$  in the lexicographic order. As we pointed out before, to find the relative order of the nodes  $u, v$  in a Wheeler graph, we need to compare the suffixes of  $T$  that precede  $u$  and  $v$ . According to the proof of Theorem 8, it only suffices to compare the suffixes  $\alpha, \beta$  of the phrases in which  $u, v$  are located, and if they are equal, we need only to compare their following suffixes, which we can do using the parse  $\mathcal{P}$ . Having already computed the Wheeler graph  $G_{\mathcal{P}}$ , we know the lexicographic rank of each suffix among those starting at phrase boundaries.

Hence, as in Boucher et al. [5], we will iterate over the  $SA_{\mathcal{D}}$  while we simultaneously build a succinct representation of  $G_T$ , i.e. the array  $L$  together with bit vectors  $I$  and  $O$ . For each phrase  $d \in \mathcal{D}$ , we save the ranks of vertices  $u \in V$  such that  $e = (u, v) \in E$  and  $\lambda(e) = d$ . Moreover, for each such phrase suffix  $s$ , we save the labels of the outgoing edges from the vertices  $v \in V$  that are suffixed by  $s$ .

<sup>2</sup> we even have the relative order of some vertices with both in-degree and out-degree equal to 1, but that is not helpful in general



We will proceed in two passes, each time outputting the entries of the arrays  $L$ ,  $I$  and  $O$  for a different subset of nodes. First, we will process those that correspond to vertices suffixed by those elements  $s \in \mathcal{S}$  that are either suffix of only a single phrase  $d \in \mathcal{D}$  or appear only once in  $\mathcal{D}$ , leaving placeholders for the rest of the vertices. The vertices corresponding to elements of  $\mathcal{S}$  that are suffixes of multiple phrases of  $\mathcal{D}$  or are a whole phrase that appears multiple times in  $T$  will be processed according to the increasing Wheeler order of their phrase-boundary suffixes. All in all, this proves Theorem 9.

► **Theorem 9.** *Let  $G_{\mathcal{P}} = (V_{\mathcal{P}}, E_{\mathcal{P}}, \lambda_{\mathcal{P}})$  be a Wheeler graph over a PFP of the string  $T$  of length  $n$ . Then we can build the expanded Wheeler graph  $G = (V, E, \lambda)$  from  $G_{\mathcal{P}}$  in  $O(n)$  time and workspace proportional to  $O(\text{PFP}(T))$ .*

### 3.2 Implementation

We have built our approach as an extension of the `bigbwt` – an implementation of the original method by Boucher et al. [5] for building BWTs using PFP. We have used the original PFP component of the `bigbwt` without changes. In contrast with `bigbwt`, the only intermediate file that we need apart from the PFP is the tunnelled BWT of the input. We have used the linear time SACA-K [18] algorithm that only uses  $O(1)$  workspace to create the suffix array of  $\mathcal{P}$ , which is subsequently transformed into the  $BWT(\mathcal{P})$ . As in `bigbwt`, the suffix array of  $\mathcal{D}$  along with the LCP array is created using the gSACAK algorithm [17] and is later used to iterate over the suffixes of the phrases from  $\mathcal{D}$  in the increasing lexicographic order.

To find the tunnelling of the  $BWT(\mathcal{P})$ , we considered several approaches suggested by Baier et al. [2, 4, 3]. Most of them aimed at reducing the compressed size of the run-length encoded BWT. Since our current goal is not to optimize the compressed size of the constructed WG (which is not straightforward to do when only having access to  $\mathcal{P}$ ), we have employed a tunnelling method which does not necessarily minimize the size of the tunnelled BWT after compression, but rather the number of edges in the resulting WG. While being tightly linked with the edge-minimization of the de Bruijn graph constructed from the input text, the prefix intervals in this method do not overlap, and Baier et al. [3] have shown that the method leads to significant reduction of BWT size for repetitive datasets. The only alteration we have made is that we have enabled the processing of large integer alphabets.

The implementation is heavily based on the Succinct Data Structure Library (SDSL) by Simon Gog et al. [12]. Namely, the tunnelled BWT is represented as a wavelet tree that stores the  $L$  component, which allows for efficient select queries. The rank and select support for the bitvectors  $I$  and  $O$  is also supported by SDSL.

Similarly to the implementation of the `bigbwt`<sup>3</sup> [5], we slightly deviate from the method of resolving the ambiguous phrase suffixes; instead of leaving the placeholders to fill in another pass, we use heapsort to merge the given vertices according to the rest of the suffixes. We note that this information is already available in the  $G_{\mathcal{P}}$  as the labels of the vertices in  $V_{\mathcal{P}}$ , so we only need to sort integers instead of actual suffixes. This approach allows us to sequentially write the succinct representation of the resulting WG in a single pass, to some extent alleviating the slow writing process of external memories.

<sup>3</sup> <https://gitlab.com/manzai/Big-BWT>

## 4 Experiments

We have demonstrated the applicability of our approach using two real-world datasets, 5520 genomes of the Salmonella genus and 1000 copies of human chromosome 19.

The Salmonella dataset was obtained using `ncbi-datasets-cli` tool. Particularly, the command `datasets download genome taxon salmonella -assembly-level complete_genome` was used. The genomic fasta files were concatenated into a single fasta file of size 13GB. The dataset of haplotypes of human chromosome 19 was obtained from the 1000 genomes project [22]<sup>4</sup>.

From these fasta files, we extracted subsets of increasing sizes to observe the trends of running time and memory consumption, as well as the sizes of the resulting WGs. For the sake of simplicity, we preprocessed the datasets so that we ignored all characters apart from A, C, G, T. All of the experiments were run on a machine with Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz and 144GB of operational memory. Every program was executed using a single thread. The running time and peak memory consumption were measured using the `/usr/bin/time -v` command.

The results of the experiments for the copies of the Salmonella genus and chromosome 19 are displayed in Table 1 and 2. Apart from our implementation, we also provide the results for the approach of Baier et al. [3] for comparison. Their approach consists of first constructing a BWT of the input and then tunnelling the constructed BWT. The BWT construction algorithm can be done by either the DIVSUF SORT algorithm [8] or a semi-external variant of the induced sorting algorithm [19] (SE-SAIS), where the former is fast but demanding significantly more memory than the latter, which is slower but memory-efficient. We have used the SE-SAIS since it makes a more fair comparison to our PFP-based algorithm. Our experiments were all ran with the PFP parameters set to  $w = 4$  and  $p = 50$ .

Expectedly, the WGs we produce are larger than those of Baier et al. due to the constraints of PFP, which limits the space of possibilities of tunnelling and treats the removal of any parse character such as having the same benefit, which is clearly not the case from the perspective of the input. We also point out that smaller WGs can be achieved with different combinations of the PFP parameters. Our approach, however, operates within much less memory and shorter running time, which allows it to construct WGs for datasets of unprecedented sizes, e.g. the whole 1000 human genomes project and beyond.

■ **Table 1** Results of the experiments on the Salmonella genomes. The 'time' is the running time in seconds, 'memory' stands for the peak memory usage in MBs during the running time. The file sizes are reported in MBs.

#sequences	size	Baier et al. [3]			PFP		
		time	memory	size	time	memory	size
1	5	8	47	6	5	57	7
10	34	42	118	16	24	223	38
100	238	301	389	55	99	696	174
1000	2619	3431	3128	430	568	2709	1400
5000	12053	18658	14538	1275	1726	4859	4189

<sup>4</sup> dataset available at <http://dolomit.cs.tu-dortmund.de/tudocomp/>

■ **Table 2** Results of the experiments on the copies of chromosome 19. The 'time' is the running time in seconds, 'memory' stands for the peak memory usage in MBs during the running time. The file sizes are reported in MBs. The Baier et al.'s [3] approach for 1000 sequences was terminated after 20 hours.

#sequences	size	Baier et al. [3]			PFP		
		time	memory	size	time	memory	size
1	54	81	341	67	82	597	70
10	533	700	708	87	157	880	296
100	5322	7263	6748	162	644	1319	2123
500	27903	54498	34031	749	2634	4190	13110
1000	53220	—	—	—	5938	8009	29803

## 5 Conclusion and future work

We have successfully demonstrated that the prefix-free parsing technique can be used to alleviate the computational requirements of the construction of tunnelled BWTs. Our approach allows the use of Wheeler graphs as pangenomic references for huge datasets such as the 1000 Genomes Project, the Vertebrate Genomes Project, the Earth Microbiome Project, and many more.

To this end, we generalized the approach of Boucher et al. [5] devised for the construction of BWTs from large volumes of repetitive data and experimentally showed it can be a good starting point for building Wheeler graphs.

Since our approach treats the PFP parse as any input text, it does not exploit the full information the PFP provides and could benefit from incorporating the lengths of the dictionary phrases into the tunnelling process, potentially even to the point of simulating the tunnelling algorithm on the original text and producing the same output. We leave this for future work.

Another logical next step is to enhance our approach with the ability to output a sufficiently small suffix array sample to allow our Wheeler graphs to locate the occurrences of patterns.

---

## References

- 1 Jarno Alanko, Giovanna D'Agostino, Alberto Policriti, and Nicola Prezza. Regular languages meet prefix sorting. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 911–930. SIAM, 2020.
- 2 Uwe Baier. On Undetected Redundancy in the Burrows-Wheeler Transform. *Annual Symposium on Combinatorial Pattern Matching (CPM 2018)*, 105:3:1–3:15, 2018. doi:10.4230/LIPIcs.CPM.2018.3.
- 3 Uwe Baier, Thomas Büchler, Enno Ohlebusch, and Pascal Weber. Edge minimization in de Bruijn graphs. In *2020 Data Compression Conference (DCC)*, pages 223–232. IEEE, 2020.
- 4 Uwe Baier and Kadir Dede. BWT Tunnel Planning is hard but manageable. In *2019 Data Compression Conference (DCC)*, pages 142–151. IEEE, 2019.
- 5 Christina Boucher, Travis Gagie, Alan Kuhnle, Ben Langmead, Giovanni Manzini, and Taher Mun. Prefix-free parsing for building big BWTs. *Algorithms for Molecular Biology*, 14(1):1–15, 2019.
- 6 Michael Burrows and David Wheeler. A block-sorting lossless data compression algorithm. In *Digital SRC Research Report*. Citeseer, 1994.

- 7 Paolo Ferragina and Giovanni Manzini. Opportunistic data structures with applications. In *Proceedings 41st annual symposium on foundations of computer science*, pages 390–398. IEEE, 2000.
- 8 Johannes Fischer and Florian Kurpicz. Dismantling divsufsort. *arXiv preprint*, 2017. [arXiv:1710.01896](https://arxiv.org/abs/1710.01896).
- 9 Travis Gagie, Giovanni Manzini, and Jouni Sirén. Wheeler graphs: A framework for BWT-based data structures. *Theoretical computer science*, 698:67–78, 2017.
- 10 Erik Garrison, Jouni Sirén, Adam M Novak, Glenn Hickey, Jordan M Eizenga, Eric T Dawson, William Jones, Shilpa Garg, Charles Markello, Michael F Lin, et al. Variation graph toolkit improves read mapping by representing genetic variation in the reference. *Nature biotechnology*, 36(9):875–879, 2018.
- 11 Daniel Gibney and Sharma V Thankachan. On the complexity of recognizing wheeler graphs. *Algorithmica*, 84(3):784–814, 2022.
- 12 Simon Gog, Timo Beller, Alistair Moffat, and Matthias Petri. From theory to practice: Plug and play with succinct data structures. In *13th International Symposium on Experimental Algorithms, (SEA 2014)*, pages 326–337, 2014. doi:10.1007/978-3-319-07959-2\_28.
- 13 Juha Kärkkäinen and Peter Sanders. Simple linear work suffix array construction. In *International colloquium on automata, languages, and programming*, pages 943–955. Springer, 2003.
- 14 Ben Langmead. Aligning short sequencing reads with Bowtie. *Current protocols in bioinformatics*, 32(1):11–7, 2010.
- 15 Ben Langmead and Steven L Salzberg. Fast gapped-read alignment with Bowtie 2. *Nature methods*, 9(4):357–359, 2012.
- 16 Heng Li. Aligning sequence reads, clone sequences and assembly contigs with BWA-MEM. *arXiv preprint*, 2013. [arXiv:1303.3997](https://arxiv.org/abs/1303.3997).
- 17 Felipe A Louza, Simon Gog, and Guilherme P Telles. Inducing enhanced suffix arrays for string collections. *Theoretical Computer Science*, 678:22–39, 2017.
- 18 Ge Nong. Practical linear-time  $O(1)$ -workspace suffix sorting for constant alphabets. *ACM Transactions on Information Systems (TOIS)*, 31(3):1–15, 2013.
- 19 Ge Nong, Sen Zhang, and Wai Hong Chan. Linear suffix array construction by almost pure induced-sorting. In *2009 data compression conference*, pages 193–202. IEEE, 2009.
- 20 Alberto Policriti and Nicola Prezza. From LZ77 to the Run-Length Encoded Burrows-Wheeler Transform, and Back. In Juha Kärkkäinen, Jakub Radoszewski, and Wojciech Rytter, editors, *28th Annual Symposium on Combinatorial Pattern Matching (CPM 2017)*, volume 78 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:10, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.CPM.2017.17.
- 21 Jarno Alanko; Nicola Cotumaccio; Nicola Prezza. Linear-time minimization of Wheeler DFAs, 2022. URL: <https://sigport.org/documents/linear-time-minimization-wheeler-dfas/>.
- 22 Nayanah Siva. 1000 genomes project. *Nature biotechnology*, 26(3):256–257, 2008.
- 23 Daniel Valenzuela and Veli Mäkinen. CHIC: a short read aligner for pan-genomic references. *bioRxiv*, page 178129, 2017.