<https://helda.helsinki.fi>

Tunneling on Wheeler Graphs

Alanko, Jarno

IEEE
2019

Alanko , J , Gagic , T , Navarro , G & Benkner , L S 2019 , Tunneling on Wheeler Graphs . in 2019 DATA COMPRESSION CONFERENCE (DCC) . IEEE Data Compression Conference , IEEE , New York , pp. 122-131 , Data Compression Conference , Snowbird , Utah , United States , 26/03/2019 . <https://doi.org/10.1109/DCC.2019.00020>

<http://hdl.handle.net/10138/308407>
<https://doi.org/10.1109/DCC.2019.00020>

cc_by
acceptedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

Tunneling on Wheeler Graphs

Jarno Alanko¹, Travis Gagie^{2,3}, Gonzalo Navarro^{2,4}, Louisa Seelbach Benkner⁵

¹Dept. of Computer Science, University of Helsinki, Finland, jarno.alanko@helsinki.fi

²CeBiB — Center for Biotechnology and Bioengineering, Chile

³EIT, Diego Portales University, Chile, travis.gagie@gmail.com

⁴Dept. of Computer Science, University of Chile, Chile, gnavarro@dcc.uchile.cl

⁵Dept. of Electrical Engineering and Computer Science, University of Siegen, Germany, seelbach@eti.uni-siegen.de

Abstract

The Burrows-Wheeler Transform (BWT) is an important technique both in data compression and in the design of compact indexing data structures. It has been generalized from single strings to collections of strings and some classes of labeled directed graphs, such as tries and de Bruijn graphs. The BWTs of repetitive datasets are often compressible using run-length compression, but recently Baier (CPM 2018) described how they could be even further compressed using an idea he called *tunneling*. In this paper we show that tunneled BWTs can still be used for indexing and extend tunneling to the BWTs of Wheeler graphs, a framework that includes all the generalizations mentioned above.

Introduction

The Burrows-Wheeler transform (BWT) is a cornerstone of data compression and succinct text indexing. It is a reversible permutation on a string that tends to compress well with run-length coding, while simultaneously facilitating pattern matching against the original string. Recently, practical data structures have been designed on top of the run-length compressed BWT to support optimal-time text indexing within space bounded by the number of runs of the BWT [1].

However, run-length coding does not necessarily exploit all the available redundancy in the BWT of a repetitive string. To this end, Baier recently introduced the concept of *tunneling* to compress the BWT by exploiting additional redundancy not yet captured by run-length compression [2]. While his representation can achieve better compression than run-length coding, no support for text indexing is given.

Meanwhile, the concept of *Wheeler graphs* was introduced by Gagie et al. as an alternative way to view Burrows-Wheeler type indices [3]. The framework can be used to derive a number of existing index structures in the Burrows-Wheeler family, like the classical FM-index [4] including its variants for multiple strings [5] and alignments [6], the XBWT for trees [7], the GCSA for directed acyclic graphs [8], and the BOSS data structure for de Bruijn graphs [9].

In this work, we show how Baier’s concept of tunneling can be neatly explained in terms of Wheeler graphs. Using the new point of view, we show how to support

FM-index style pattern searching on the tunneled BWT. We also describe a sampling strategy to support pattern counting and locating and character extraction, making our set of data structures a fully-functional FM-index. Supporting FM-index operations was posed as an open problem by Baier [2].

We also use the generality of the Wheeler graph framework to generalize the concept of tunneling to any Wheeler graph. This result can be used to compress any Wheeler graph while still supporting basic pattern searching to decide if a pattern exists as a path label in the graph. This has applications for all index structures that can be explained in terms of Wheeler graphs.

Preliminaries

Let $\mathcal{G} = (V, E, \lambda)$ denote a directed edge-labeled multi-graph, in which V denotes the set of nodes, E denotes the multiset of edges and $\lambda : E \rightarrow A$ denotes a function labeling each edge of $\mathcal{G}$ with a character from a totally-ordered alphabet A . Let $\prec$ denote the ordering among A 's elements. We follow the definition of Gagie et al. [3].

Definition 1 (Wheeler graph). The graph $\mathcal{G} = (V, E, \lambda)$ is called a *Wheeler graph* if there is an ordering on the set of nodes such that nodes with in-degree 0 precede those with positive in-degree and for any two edges $(u_1, v_1), (u_2, v_2)$ labeled with $\lambda((u_1, v_1)) = a_1$ and $\lambda((u_2, v_2)) = a_2$, we have

- (i) $a_1 \prec a_2 \Rightarrow v_1 < v_2$,
- (ii) $(a_1 = a_2) \wedge (u_1 < u_2) \Rightarrow v_1 \leq v_2$. □

We note that the definition implies that all edges arriving at a node have the same label. We call such an ordering a *Wheeler ordering of nodes*, and the rank of a node within this ordering the *Wheeler rank* of the node. Given a pattern $P \in A^*$, we call a node $v \in V$ an *occurrence* of P if there is a path in $\mathcal{G}$ ending at v such that the concatenation of edge labels on the path is equal to P . The Wheeler ranks of all occurrences of P form a contiguous range of integers, which we call the *Wheeler range* of P . Finding such a range is called *path searching* for P .

Given a Wheeler ordering of nodes, we define the corresponding *Wheeler ordering of edges* such that for a pair of edges $e_1 = (u_1, v_1)$ and $e_2 = (u_2, v_2)$, we have $e_1 < e_2$ iff $\lambda(e_1) \prec \lambda(e_2)$ or $(\lambda(e_1) = \lambda(e_2) \text{ and } u_1 < u_2)$. When referring to edges, the term *Wheeler rank* refers to the rank of the edge in the Wheeler ordering of edges, and *Wheeler range* refers to a set of edges whose Wheeler ranks form a contiguous interval.

We use a slightly modified version of the representation of Wheeler graphs proposed by Gagie et al. [3]. Suppose we have a Wheeler graph with n nodes and m edges. Then we represent the graph with the following data structures:

- A string $L[1..m] = L_1 \cdots L_n$ where L_i is the concatenation of the labels of the $|L_i|$ edges going out from the node with Wheeler rank i such that the labels from a node are concatenated in their relative Wheeler order.
- An array $C[1..|A|]$ such that $C[i]$ is the number of edges in E with a label smaller than the i th smallest symbol in A .

- A binary string $I[1..n+m+1] = X_1, \dots, X_n \cdot 1$ where $X_i = 1 \cdot 0^{k_i}$ and k_i is the indegree of the node with Wheeler rank i .
- A binary string $O[1..n+m+1] = X_1, \dots, X_n \cdot 1$ where $X_i = 1 \cdot 0^{l_i}$ and l_i is the outdegree of the node with Wheeler rank i .

Given these data structures, we can traverse the Wheeler graph with the following two operations: First, given the Wheeler rank i of a node, find the Wheeler rank of the k -th out-edge labeled c from the node using Eq. (1) below; second, given the Wheeler rank j of an edge, find the Wheeler-rank r of its target node with Eq. (2):

$$C[c] + \text{rank}_c(L, \text{select}_1(O, i) - i) + k, \quad (1)$$

$$\text{rank}_1(I, \text{select}_0(I, j)), \quad (2)$$

where, given a string S , $\text{rank}_c(S, i)$ denotes the number of occurrences of character c in $S[1..i]$ and $\text{select}_c(S, i)$ denotes the position of the i -th occurrence of c in S . For the binary strings I and O , these operations can be carried out in constant time using $o(|I| + |O|)$ extra bits [10]. For string L , operation rank can be carried out in time $O(\log \log_w |A|)$ on a w -bit RAM machine using $o(|L| \log |A|)$ extra bits [11, Thm. 8].

Given a Wheeler range $[i, i']$ of nodes, we can find the Wheeler rank of the first and last edge labeled c leaving from a node in $[i, i']$ with a variant of Eq. (1): $C[c] + \text{rank}_c(L, \text{select}_1(O, i) - i) + 1$ and $C[c] + \text{rank}_c(L, \text{select}_1(O, i' + 1) - (i' + 1))$, respectively; then we apply Eq. (2) on both edge ranks to obtain the corresponding node range (the result is a range by the path coherence property of Wheeler graphs [3]). This operation enables path searches on Wheeler graphs.

Tunneling

We adapt Baier's concept of *blocks* on the BWT [2] to Wheeler graphs.

Definition 2 (Block). A *block* $\mathcal{B}$ of a Wheeler graph $\mathcal{G} = (V, E, \lambda)$ of size s and width w is a sequence of w -tuples $(v_{1,1}, \dots, v_{w,1}), \dots, (v_{1,s}, \dots, v_{w,s})$ of pairwise distinct nodes of $\mathcal{G}$ such that

- (i) For $1 \leq i \leq w - 1$ and $1 \leq j \leq s$, the node $v_{i+1,j}$ is the immediate successor of $v_{i,j}$ with respect to the Wheeler ordering on V .
- (ii) For $1 \leq i \leq w$, let $V_i = \{v_{i,j} \mid 1 \leq j \leq s\}$, $E_i = E \cap (V_i \times V_i)$, and $\lambda_i = \lambda|_{E_i}$. The subgraphs $t_i = (V_i, E_i, \lambda_i)$ are isomorphic subtrees of $\mathcal{G}$, preserving topology and labels. For $1 \leq i \leq w - 1$, let $f_i : t_i \rightarrow t_{i+1}$ denote the corresponding isomorphisms, thus $v_{i+1,j} = f_i(v_{i,j})$ for all $1 \leq j \leq s$.
- (iii) For $1 \leq i \leq w$, let $v_{i,1}$ denote the root node of t_i . In particular, $v_{i,1}$ is the only node of indegree 0 in t_i . All edges leading to a node in $\{v_{i,1} \mid 1 \leq i \leq w\}$ are labeled with the same character. The indegrees of these nodes may differ.
- (iv) For $1 \leq i \leq w$ and $2 \leq j \leq s$, the nodes $v_{i,j}$ are of indegree 1 in $\mathcal{G}$ (and by (i) and (ii), of indegree 1 in the corresponding subtree t_i , that is, the only edge in $\mathcal{G}$ leading to such a node $v_{i,j}$ belongs to the subtree t_i).

Figure 1: Tunneling a block of size 7 and width 2 in a Wheeler graph.

- (v) For every integer $1 \leq j \leq s$ and character $c \in A$, exactly one of the following conditions holds:
 - (a) For every $1 \leq i \leq w$, there is exactly one out-edge of $v_{i,j}$ labeled with c , which is contained in E_i . There are no out-edges of $v_{i,j}$ labeled with c leading to non-block nodes.
 - (b) For every $1 \leq i \leq w$, there is no out-edge of $v_{i,j}$ labeled with c contained in E_i . There may be out-edges of $v_{i,j}$ labeled with c leading to non-block nodes. The number of such out-edges for each node may differ. $\square$

A block $\mathcal{B} = ((v_{1,1}, \dots, v_{w,1}), \dots, (v_{1,s}, \dots, v_{w,s}))$, abbreviated $\mathcal{B} = (v_{i,j})_{1 \leq i \leq w, 1 \leq j \leq s}$, is called *maximal in size* if, for any choice of nodes $v_1, \dots, v_w \in V$, the sequences of size $s + 1$ of w -tuples $((v_i)_{1 \leq i \leq w}, (v_{i,1})_{1 \leq i \leq w}, \dots, (v_{i,s})_{1 \leq i \leq w})$ and $((v_{i,1})_{1 \leq i \leq w}, \dots, (v_{i,s})_{1 \leq i \leq w}, (v_i)_{1 \leq i \leq w})$ do not form a block. The block is called *maximal in width* if, for any choice of nodes $v_1, \dots, v_s \in V$, the sequences of $(w+1)$ -tuples $((v_j, v_{1,j}, \dots, v_{w,j}))_{1 \leq j \leq s}$ and $((v_{1,j}, \dots, v_{w,j}, v_j))_{1 \leq j \leq s}$ do not form a block. A block is called *maximal* if it is maximal in both width and size.

Let $\mathcal{G} = (V, E, \lambda)$ denote a Wheeler graph containing a maximal block. We then obtain a directed edge-labeled (multi-)graph $\mathcal{G}_t$ from $\mathcal{G}$ as follows:

- (i) We merge the corresponding nodes and edges of the isomorphic subtrees t_i , with $1 \leq i \leq w$, in order to obtain a subtree t of $\mathcal{G}$. In particular, for $1 \leq j \leq s$, we collapse the nodes of the w -tuple $(v_{1,j}, \dots, v_{w,j})$ to obtain a node x_j of the graph t . The labels of the merged edges coincide and stay the same.
- (ii) All edges leading from a non-block node u to the root node $v_{i,1}$ of the subtree t_i are redirected to lead to the node x_1 of t , preserving their labels.

- (iii) All edges leading from a node $v_{i,j}$ of subgraph t_i to a non-block node u are redirected to leave from the node x_j of t , preserving their labels.

Formally, the graph $\mathcal{G}_t = (V_t, E_t, \lambda_t)$ is defined as follows: The set of nodes of $\mathcal{G}_t$ is defined as $V_t = (V \setminus \{v_{i,j} \mid 1 \leq i \leq w, 1 \leq j \leq s\}) \cup \{x_j \mid 1 \leq j \leq s\}$. We define a function $\varphi : V \rightarrow V_t$ mapping a node in $\mathcal{G}$ to its corresponding node in $\mathcal{G}_t$ by

$$\varphi(v) = \begin{cases} v & \text{if } v \notin \{v_{i,j} \mid 1 \leq i \leq w, 1 \leq j \leq s\} \\ x_j & \text{if } v = v_{i,j} \text{ for some integers } 1 \leq i \leq w, 1 \leq j \leq s. \end{cases}$$

The multiset of edges E_t is defined as the difference of multisets $\{(\varphi(u), \varphi(v)) \mid (u, v) \in E\} \setminus \{(\varphi(u), \varphi(v)) \mid (u, v) \in E_i, 2 \leq i \leq w\}$. For every edge $(x, y) \in E_t$, there is a corresponding edge $(u, v) \in E$ such that $(x, y) = (\varphi(u), \varphi(v))$. Thus, we define $\lambda_t((x, y)) = \lambda((u, v))$. This is well-defined because only edges with the same label are merged. We call *tunneling* the process of obtaining the graph $\mathcal{G}_t$ from $\mathcal{G}$. See Fig. 1.

Lemma 3. *Let $\mathcal{G} = (V, E, \lambda)$ denote a Wheeler graph containing a maximal block $\mathcal{B} = (v_{i,j})_{1 \leq i \leq w, 1 \leq j \leq s}$ of width w and size s and let $\mathcal{G}_t = (V_t, E_t, \lambda_t)$ denote the graph obtained from $\mathcal{G}$ by tunneling. Then $\mathcal{G}_t$ is a Wheeler graph.*

Proof. To show that $\mathcal{G}_t$ is a Wheeler graph, we must define a Wheeler ordering on V_t . As only consecutive nodes of V are merged in order to obtain $\mathcal{G}_t$, this induces a canonical ordering on the nodes of V_t : Pick two nodes $x \neq y$ of V_t . Let $\varphi^{-1}(x), \varphi^{-1}(y) \subseteq V$ denote the corresponding preimages under φ . By property (i) of Def. 2, the nodes of $\varphi^{-1}(x)$ (respectively, $\varphi^{-1}(y)$) are consecutive in Wheeler order. Thus, we either have $u < v$ for every $u \in \varphi^{-1}(x), v \in \varphi^{-1}(y)$, or vice versa. We set $x < y$ in the first case and $x > y$ in the second. This yields an ordering on the nodes of V_t such that for any two nodes $u \neq v$ of V , we have $\varphi(u) < \varphi(v) \Rightarrow u < v$ and $u < v \Rightarrow \varphi(u) \leq \varphi(v)$.

First, take two nodes $x \neq y$ of $\mathcal{G}_t$, such that x has in-degree 0 and y is of positive in-degree. As in the process of tunneling the in-degree of a node is not decreased, every node $u \in \varphi^{-1}(x)$ is of in-degree 0 as well. Moreover, as y is of positive in-degree, there is a node $v \in \varphi^{-1}(y)$, such that v is of positive in-degree. As $\mathcal{G}$ is a Wheeler graph, we have $u < v$ and thus $x = \varphi(u) \leq \varphi(v) = y$. As by assumption, $x \neq y$, this yields $x < y$. Therefore, nodes with in-degree 0 precede those with positive in-degree.

Now, take two edges (x, y) and (x', y') of $\mathcal{G}$ labeled with a and a' , respectively. Without loss of generality, assume $a \preceq a'$. Choose $u \in \varphi^{-1}(x), v \in \varphi^{-1}(y), u' \in \varphi^{-1}(x')$ and $v' \in \varphi^{-1}(y')$, such that (u, v) and (u', v') are edges of $\mathcal{G}$. By definition of $\mathcal{G}_t$, the label on the edge (u, v) (resp., (u', v')) is a (resp., a'). We then have $(x, y) = (\varphi(u), \varphi(v))$ and $(x', y') = (\varphi(u'), \varphi(v'))$. Consider the two cases of Def. 1:

- (i) Let $a \prec a'$. As $\mathcal{G}$ is a Wheeler graph, we have $v < v'$ and thus $\varphi(v) \leq \varphi(v')$. If $\varphi(v) < \varphi(v')$ we are done, so assume $\varphi(v) = \varphi(v')$. We then have $v = v_{i,j}$ and $v' = v_{k,j}$ for some nodes $v_{i,j} \neq v_{k,j}$ of the block. By properties (ii) - (iv) of Def. 2, the labels of the incoming edges of $v_{i,j}$ and $v_{k,j}$ are the same, contradicting $a \prec a'$.
- (ii) Let $a = a'$ and without loss of generality assume $u < u'$. This yields $\varphi(u) \leq \varphi(u')$. As $\mathcal{G}$ is a Wheeler graph, we obtain $v \leq v'$, and thus $\varphi(v) \leq \varphi(v')$. $\square$

Two blocks $\mathcal{B}_1 = (v_{i,j})_{1 \leq i \leq w_1, 1 \leq j \leq s_1}$ and $\mathcal{B}_2 = (u_{i,j})_{1 \leq i \leq w_2, 1 \leq j \leq s_2}$ of a Wheeler graph $\mathcal{G} = (V, E, \lambda)$ are called *disjoint* if their corresponding node sets $\{v_{i,j} \mid 1 \leq i \leq w_1, 1 \leq j \leq s_1\} \subset V$ and $\{u_{i,j} \mid 1 \leq i \leq w_2, 1 \leq j \leq s_2\} \subset V$ are disjoint. Since, by Lemma 3, a graph obtained from a Wheeler graph by tunneling is still a Wheeler graph, we can tunnel iteratively with disjoint blocks.

Definition 4 (Tunneled Graph). Let $\mathcal{G}$ be a Wheeler graph containing k pairwise disjoint maximal blocks $\mathcal{B}_1, \dots, \mathcal{B}_k$. The *tunneled graph* $\mathcal{G}_t$ of $\mathcal{G}$ corresponding to those blocks is defined as the Wheeler graph obtained from $\mathcal{G}$ by iteratively tunneling all the blocks $\mathcal{B}_i$, for $1 \leq i \leq k$. Each maximal block $\mathcal{B}_i$ is also called a *tunnel*. $\square$

Note that tunnels more general than the tree form given in Def. 2 would break the Wheeler graph rules of Def. 1 before or after tunneling, except that the nodes of t_1 and t_w could be connected with outside nodes, all of Wheeler ranks smaller and larger, respectively, than their corresponding tunnel edges. We could also handle forests, but those can be seen as a set of disjoint tunnels.

Path searching on a tunneled Wheeler graph

Wheeler graphs can be searched for the existence of paths whose concatenated labels yield a given string P [3], generalizing the classical backward search on strings [4]. We now show that those searches can also be performed on tunneled Wheeler graphs.

Given $\mathcal{G}_t$, we can simulate the traversal of $\mathcal{G}$ as follows. A node $v \in V$ is represented by a pair $\langle \varphi(v), \text{off}(v) \rangle$, where $\varphi(v)$ is the corresponding node in $\mathcal{G}_t$ as defined in the previous section and $\text{off}(v)$ is the *tunnel offset* of v . If the node $\varphi(v)$ does not belong to a tunnel, then it must be that $\text{off}(v) = 1$ and the pair represents just the node v . Otherwise, $\varphi(v)$ corresponds to multiple nodes $(v_{i,j})_{1 \leq i \leq w}$ of some tunnel $\mathcal{B} = (v_{i,j})_{1 \leq i \leq w, 1 \leq j \leq s}$ in $\mathcal{G}$ and the tunnel offset represents which of the original nodes we are currently at, in Wheeler rank order. That is, if $v = v_{i,j}$, then $\text{off}(v) = i$.

The idea is that, when our traversal enters a tunnel $\mathcal{B}$, we remember which original subgraph t_i we actually entered, and use that information to exit the tunnel accordingly. We mark the nodes of $\mathcal{G}_t$ that are tunnel entrances in a bitvector, and the other tunnel nodes in another bitvector. We then distinguish three cases.

Keeping out of tunnels. If we are at a pair $\langle i, 1 \rangle$ not in a tunnel, compute j and r with Eqs. (1) and (2), respectively, and it turns out that r is not marked as a tunnel entrance, then we stay out of any tunnel and our new pair is $\langle r, 1 \rangle$.

Entering a tunnel. Assume we are at a pair $\langle i, 1 \rangle$, where $i = \varphi(v)$ for some non-tunnel node v , compute j and r with Eqs. (1) and (2), respectively, and then it turns out that $r = \varphi(u)$ is marked as a tunnel entrance. Then we have entered a tunnel and the new pair must be $\langle r, o \rangle$, for some offset o we have to find out.

The w nodes $(u_p)_{1 \leq p \leq w} \in V$ that were collapsed to form $\varphi(u)$ are of indegree 0 within the subgraphs t_p , but may receive a number of edges from non-tunnel nodes (indeed, we are traversing one). Since all those edges are labeled by the same symbol

c , the Wheeler rank of all the sources of edges that lead to u_p must precede the Wheeler ranks of all the sources of edges that lead to $u_{p'}$ for any $1 \leq p < p' \leq w$.

The problem is, knowing that we are entering by the edge with Wheeler rank j , and that the edges that enter into r start at Wheeler rank $\text{select}_1(I, r) - r + 1$, how to determine the index o of the subgraph t_o we have entered. For this purpose, we store a bitvector $I'[1..m]$, where $m = |E_t|$, so that $I'[j] = 1$ iff the j th edge of E_t , in Wheeler order, corresponds to the first edge leading to its target in $\mathcal{G}$. Said another way, I' marks, in the area of I corresponding to the edges that reach $\varphi(u)$, which were the first edges arriving at each copy u_p that was collapsed to form $\varphi(u)$. We can then compute $o = \text{rank}_1(I', j) - \text{rank}_1(I', \text{select}_1(I, r) - r)$.

Moving in a tunnel. Assume we are at a pair $\langle i, o \rangle$ inside a tunnel, for $i = \varphi(u)$, and want to traverse the k th edge labeled c leaving the pair. We then use the formulas given after Eqs. (1) and (2) to compute the first and last edge labeled c leaving node i . These form a Wheeler range $[j_1, j_2]$. We then apply Eq. (2) from $j = j_1$ to find the first target node r . If r is marked as an in-tunnel node, then we know that letter c keeps us inside the tunnel, $j_1 = j_2$ (that is, there is exactly one out-edge by letter c from each of the nodes $(u_p)_{1 \leq p \leq w}$ that were collapsed to form $\varphi(u)$, by part (a) of item (v) of Definition 2), and thus our new node is simply $\langle r, o \rangle$.

If, instead, r is not an in-tunnel node, then letter c takes us out of the tunnel, and we must compute the appropriate target node. Each of the nodes u_s may have zero or more outgoing edges labeled c . The Wheeler ranks of the edges leaving u_p precede those of the edges leaving $u_{p'}$, for any $1 \leq p < p' \leq w$. Thus, we use a bitvector $O'[1..m]$ analogous to I' , where $O'[j] = 1$ iff the j th edge of E_t in Wheeler order corresponds to the first edge leaving a node in $\mathcal{G}$ by some letter c . In the range $O'[j_1, j_2]$, then, each 1 marks the first edge leaving from each u_p . The k th edge labeled c leaving from u_o is thus $j = \text{select}_1(O', \text{rank}_1(O', j_1) + o - 1) + k - 1$. If, for pattern searching, we want the last edge labeled c leaving from u_o , this is $j = \text{select}_1(O', \text{rank}_1(O', j_1) + o) - 1$. We then compute the correct target node rank r using Eq. (2).

Finally, there are two possibilities. If the new node r is marked as a tunnel entrance, then we have left our original tunnel to enter a new one. We then apply the method described to enter a tunnel from the values j and r we have just computed. Otherwise, r is not a tunnel node and we just return the pair $\langle r', 1 \rangle$.

Therefore, we can simulate path searching in $\mathcal{G}$ by using just our representation of $\mathcal{G}_t$. Given a character c and the Wheeler range $[i, i']$ of a string S , we can find the Wheeler range of the string Sc by following the first and last edge labeled with c leaving from $[i, i']$, as described after Eqs. (1) and (2), and operate as described from the corresponding ranks j and r . Thus, after $|P|$ steps, we have the Wheeler range of the nodes that can be reached by following the characters in P .

Theorem 5. *We can represent a Wheeler graph $\mathcal{G}$ with labeled edges from the alphabet $[1..\sigma]$ in $n_t \log \sigma + o(n_t \log \sigma) + O(n_t)$ bits of space, where n_t is the number of edges in a tunneled version of $\mathcal{G}$, such that we can decide if there exists a path labeled with P in time $O(|P| \log \log_w \sigma)$ in a w -bit RAM machine.*

Wheeler graphs of strings

We now focus on a particular type of Wheeler graphs, which corresponds to the traditional notion of Burrows-Wheeler Transform (BWT) [12] and FM-index [4] (only that our arrows go forward in the text, not backwards), and show that the full self-index functionality on strings can still be supported after tunneling. This is close to the original tunneling concept developed by Baier [2], to which we now add search and traversal capabilities.

Definition 6 (Wheeler graph of a string). Let T be a string over an alphabet A . The *Wheeler graph of the string T* is defined as $\mathcal{G} = (V, E, \lambda)$ with $V := \{v_1, \dots, v_{|T|+1}\}$, $E = \{(v_i, v_{i+1}) \mid 1 \leq i \leq |T| + 1\}$ and $\lambda : E \rightarrow A$ with $\lambda((v_i, v_{i+1})) = T[i]$, where $T[i]$ denotes the i th character in the string T . $\square$

In other words, the Wheeler graph of a string T is a path of length $|T| + 1$, where the i th edge is labeled with the i th character of T . There is exactly one valid Wheeler-ordering, which is given by the colexicographic order of prefixes of T , i.e., node v_i comes before node v_j iff the reverse of $T[1..i - 1]$ is lexicographically smaller than the reverse of $T[1..j - 1]$. There is a close connection to the BWT: the Wheeler order is given by the suffix array of the reverse of T and therefore the L -array corresponds to the BWT of the reverse of T .

For this special case of Wheeler graphs, Def. 2 simplifies as follows: A block $\mathcal{B}$ in $\mathcal{G}$ of width w and length s is a sequence of length $s + 1$ of w -tuples $(v_{1,1}, \dots, v_{w,1}), \dots, (v_{1,s+1}, \dots, v_{w,s+1})$ of pairwise distinct nodes of $\mathcal{G}$ satisfying

- (i) For $1 \leq i \leq w - 1$ and $1 \leq j \leq s + 1$, the immediate successor of the node $v_{i,j}$ with respect to the Wheeler ordering on V is $v_{i+1,j}$.
- (ii) For $1 \leq i \leq w$ and $1 \leq j \leq s$, $(v_{i,j}, v_{i,j+1})$ is an edge of E .
- (iii) For $1 \leq j \leq s$, all the edges leading to the nodes in $\{v_{i,j} \mid 1 \leq i \leq w\}$ have the same label.

The process of tunneling in a Wheeler graph $\mathcal{G}$ of a string then consists of collapsing the nodes of each w -tuple $(v_{1,j}, \dots, v_{w,j})$ into a single node x_j and collapsing the edges in $\{(v_{i,j}, v_{i,j+1}) \mid 1 \leq i \leq w\}$ into a single edge (x_j, x_{j+1}) . Furthermore, all edges leading to a node $v_{i,1}$ for some integer $1 \leq i \leq w$ are redirected to lead to the node x_1 and all edges leaving from a node $v_{i,s}$ for some integer $1 \leq i \leq w$ are redirected to leave from the node x_s . The labels of the edges stay the same. Note that we ensure that every path of the tunnel is followed by a non-tunnel node.

Suppose we have the Wheeler graph $\mathcal{G} = (V, E, \lambda)$ of a string T . Denote $|T| = n$. Let $\mathcal{G}_t = (V_t, E_t, \lambda_t)$ be a tunneled version of $\mathcal{G}$ with $|V_t| = n_t$. We represent $\mathcal{G}_t$ with the data structures L , C , I and O described in the preliminaries. We can do path searches without the bitvectors I' and O' used in the previous section, as in these particular graphs they are all 1s. We now describe how to implement the operations count, locate and extract, analogous to the operations in a regular FM-index, by using sampling schemes that extend those of the standard FM-index solution.

First, for each tunnel of length at least $\log n_t$ in $\mathcal{G}_t$, we store a pointer and the distance to the end of the tunnel for every $(\log n_t)$ th consecutive node in the tunnel. This information takes $O(n_t)$ bits of space and lets us skip to the end of the tunnel in $O(\log n_t)$ steps by walking forward until the end of the tunnel is found, or until we hit a node with a pointer to the end. The stored distance value tells us how many nodes we have skipped over.

Locating. We define a graph $\mathcal{G}_c$, called the *contracted graph*, that is identical to $\mathcal{G}_t$ except that tunnels have been contracted into single nodes. Let $\psi : \mathcal{G}_t \rightarrow \mathcal{G}_c$ be the mapping such that nodes in a tunnel in $\mathcal{G}_t$ map to the corresponding contracted node in $\mathcal{G}_c$. Take the path $Q = (q_1, \dots, q_n)$ of all nodes in $\mathcal{G}$ in the order of the path from the source to the sink. Let $Q_c = ((\psi \circ \varphi)(q_1), \dots, (\psi \circ \varphi)(q_n))$ be the corresponding path in $\mathcal{G}_c$. Let Q'_c be the same sequence as Q_c except that every run of the same node is contracted to length 1. Note that this path traverses all edges of $\mathcal{G}_c$ exactly once, i.e., it is an Eulerian path. We store a sample for every $(\log n)$ th node in the path on Q'_c , except that if a node represents a contracted tunnel, we sample the next node (our definition of blocks guarantees that the next node is not in a tunnel). The value associated with the sample is the text position corresponding to the node. This takes space $O(n_c)$, where n_c is the number of nodes in $\mathcal{G}_c$.

We can then locate the text position of a node by walking to the next sample in text order, using at most $\log n$ graph traversal operations in $\mathcal{G}_t$. In this walk we may have to skip tunnels, which is done in $O(\log n_t)$ time using their stored pointers when necessary. In the end, we subtract the travelled distance from the text position of the sampled node to get the text position of the original node. We can view such a search as a walk in $\mathcal{G}_c$, where traversing a contracted-tunnel node takes $O(\log n_t)$ graph traversal steps and traversing a non-tunneled node takes just one. The worst case time, dominated by the time to traverse tunnels, is $O(\log n \log n_t)$ steps.

Counting. Efficiently counting the number of occurrences of a pattern given its Wheeler-range requires a sampling structure different from that used for locating. The Wheeler-range could span many tunnels, whose widths are not immediately available. Let us define $w(v)$ as the width of the tunnel v belongs to, or 1 if v does not belong to a tunnel. We can then afford to sample the cumulative sum of the values $w(v)$ for all the nodes v up Wheeler-rank k for every k multiple of $\log n_t$ nodes, using $O(n_t)$ bits of space. Within this space we can also mark which nodes belong to tunnels.

This allows us to compute the sum of values $w(v)$ for any Wheeler range with endpoints that are multiples of $\log n_t$, which leaves us to compute the width of only $O(\log n_t)$ nodes at the ends of the range. For these nodes, we add 1 if they are not in a tunnel; otherwise we go to the end of the tunnel using the stored pointers and compute the width of the tunnel by looking at the out-degree of the exit of the tunnel. The total counting time is then $O(\log^2 n_t)$ graph traversal steps.

Extracting. To extract characters from T , we use a copy of the samples $\langle \text{Wheeler rank of graph node } v, \text{ text position of node } v \rangle$ we store for locating, but sorted by

text position. Also, at the end of every tunnel of length at least $\log n_t$, we store backpointers to the nodes storing pointers to the end of the tunnel.

Suppose we want to extract $T[i..j]$. If we know the node u of $\mathcal{G}_t$ representing position i , we can simply walk forward from that node to find the $j - i + 1$ desired characters by accessing L at each position. Therefore it is enough to show how to find the node u . We binary search our sample pairs to find the Wheeler rank of the closest sample before text position i . This sample is at most $\log n$ nodes away (in $\mathcal{G}_c$) from u , so we can reach u in $O(\log n)$ steps in $\mathcal{G}_c$, or equivalently, $O(\log n \log n_t)$ steps in $\mathcal{G}_t$. Note, however, that our target node u might be in a tunnel. If the tunnel is of length less than $\log n_t$, we walk towards it normally. If u is inside a tunnel of length at least $\log n_t$, instead, we use its pointers to skip to the end of the tunnel, and from there take the backpointer to the nearest position before u ; we then walk the (at most) $\log n_t$ nodes until reaching u .

The time needed to reach u is again dominated by the time to skip over and within tunnels, so the total time complexity is $O(\log n \log n_t)$ graph traversal steps.

We note that, in all cases, our graph traversal steps are of a particular form, because all the edges leaving from the current node are labeled by the same symbol. That is, the c in Eq. (1) is always $L[\text{select}_1(O, i)]$. This particular form of *rank* is called *partial rank* and it can be implemented in constant time using $o(|L| \log |A|)$ further bits [13, Lem. 2]. The following theorem summarizes the results in this section.

Theorem 7. *We can store a text $T[1..n]$ over alphabet $[1..\sigma]$ in $n_t \log \sigma + o(n_t \log \sigma) + O(n_t)$ bits of space, such that in a w -bit RAM machine we can decide the existence of any pattern P in time $O(|P| \log \log_w \sigma)$, and then report the text position of any occurrence in time $O(\log n \log n_t)$ or count the number of occurrences in time $O(\log^2 n_t)$. We can also extract any k consecutive characters of T in time $O(k + \log n \log n_t)$, where $n_t \leq n$ is the number of nodes in a tunneled Wheeler graph of T .*

Future work

Open problems are: How to find the optimal blocks that minimize space? Can we still support path searching if blocks are overlapping? Can the $O(\log^2 n)$ times of counting, locating, and extracting be reduced to $O(\log n)$, as in the basic sampling scheme on non-tunneled BWTs? How to extend those operations to more complex graphs, like trees? And can we count paths instead of path endpoints?

Acknowledgements. Funded in part by EU’s Horizon 2020 research and innovation programme under Marie Skłodowska-Curie grant agreement No 690941 (project BIRDS). T.G. and G.N. partially funded with Basal Funds FB0001, Conicyt, Chile. T.G. partly funded by Fondecyt grant 1171058. L.S.B. supported by the EU project 731143 - CID and the DFG project LO748/10-1 (QUANT-KOMP). Operation extract was designed in StringMasters 2018. We thank Uwe Baier for helpful discussions.

References

- [1] T. Gagie, G. Navarro, and N. Prezza, “Optimal-time text indexing in BWT-runs bounded space,” in *Proc. 29th SODA*, 2018, pp. 1459–1477.
- [2] U. Baier, “On undetected redundancy in the Burrows-Wheeler transform,” in *Proc. 29th CPM*, 2018, pp. 3.1–3.15.
- [3] T. Gagie, G. Manzini, and J. Sirén, “Wheeler graphs: A framework for BWT-based data structures,” *Theoretical Computer Science*, vol. 698, pp. 67–78, 2017.
- [4] P. Ferragina and G. Manzini, “Indexing compressed texts,” *Journal of the ACM*, vol. 52, no. 4, pp. 552–581, 2005.
- [5] S. Mantaci, A. Restivo, G. Rosone, and M. Sciortino, “An extension of the Burrows–Wheeler transform,” *Theoretical Computer Science*, vol. 387, no. 3, pp. 298–312, 2007.
- [6] J. C. Na, H. Kim, H. Park, T. Lecroq, M. Léonard, L. Mouchard, and K. Park, “FM-index of alignment: A compressed index for similar strings,” *Theoretical Computer Science*, vol. 638, pp. 159–170, 2016.
- [7] P. Ferragina, F. Luccio, G. Manzini, and S. Muthukrishnan, “Structuring labeled trees for optimal succinctness, and beyond,” in *Proc. 46th FOCS*, 2005, pp. 184–193.
- [8] J. Sirén, N. Välimäki, and V. Mäkinen, “Indexing graphs for path queries with applications in genome research,” *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol. 11, no. 2, pp. 375–388, 2014.
- [9] A. Bowe, T. Onodera, K. Sadakane, and T. Shibuya, “Succinct de bruijn graphs,” in *Proc. WABI*, 2012, pp. 225–235.
- [10] D. R. Clark, *Compact PAT Trees*, Ph.D. thesis, University of Waterloo, Canada, 1996.
- [11] D. Belazzougui and G. Navarro, “Optimal lower and upper bounds for representing sequences,” *ACM Transactions on Algorithms*, vol. 11, no. 4, pp. article 31, 2015.
- [12] M. Burrows and D. Wheeler, “A block sorting lossless data compression algorithm,” Tech. Rep. 124, Digital Equipment Corporation, 1994.
- [13] D. Belazzougui and G. Navarro, “Alphabet-independent compressed text indexing,” *ACM Transactions on Algorithms*, vol. 10, no. 4, pp. article 23, 2014.