Fast Exact String to D-Texts Alignments

Njagi Moses Mwaniki
Department of Computer Science, University of Pisa, Italy

Erik Garrison
University of Tennessee Health Science Center, USA

Nadia Pisanti
Department of Computer Science, University of Pisa, Italy

Abstract

In recent years, aligning a sequence to a pangenome has become a central problem in genomics and pangenomics. A fast and accurate solution to this problem can serve as a toolkit to many crucial tasks such as read-correction, Multiple Sequences Alignment (MSA), genome assemblies, variant calling, just to name a few. In this paper we propose a new, fast and exact method to align a string to a D-string, the latter possibly representing an MSA, a pan-genome or a partial assembly.
An implementation of our tool dsa is publicly available at `https://github.com/urbanslug/dsa`.

2012 ACM Subject Classification 68W01, 68W32, 68Q25, 68Q17

Keywords and phrases D-texts, Pangenomes, Alignments

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.XX

Funding
This paper is part of a project that has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 956229. Co-financed by the Connecting Europe Facility of the European Union.

arXiv:2206.03242v1 [cs.DS] 7 Jun 2022

1 Introduction

In recent years, aligning a sequence to a pangenome has become a central problem in genomics and pangenomics ([5]). This problem has often been addressed with the purpose of aligning sequencing reads to a complex structure such as a degenerate string or a more general genome graph structure (e.g.[9, 22, 2, 30, 29, 19]), and hence assuming the query string to be considerably shorter than the pangenome (e.g. [9, 22, 3, 1, 2]). As sequencing data is recently offering longer sequences with high accuracy, the attention has very recently been moved towards the problem of aligning to a graph-like structure a string which is approximately as long: a fast and accurate solution to this problem would serve as a toolkit to many crucial tasks such as read-correction, Multiple Sequences Alignment (MSA), genome assemblies, variant calling, just to name a few. To this purpose, because the requirement is an alignment that is at the same time exact and sensible to indels as well as mismatches, then new computational challenges are faced: current exact methods that perform exact alignments to long strings would have a quadratic complexity with high memory requirements, and hence be prohibitively slow.
Consider the following MSA of three closely-related sequences (on the left) and its compact representation as a D-string (on the right):

```
GCAATCGGGTATT
GCAATCGGGAATT
GCACGCTGGATTT
```

$$\hat{T} = \texttt{GCA}[\texttt{AT/CG}]\texttt{C}\,[\texttt{G/T}]\,\texttt{GG}[\texttt{TA/AA/AT}]\texttt{TT}$$

The D-string (or D-text) $\hat{T}$ contains some deterministic (shown in black) and some non-deterministic (shown in red) segments. Formally, $\hat{T}$ is a sequence of n sets of strings

© Njagi Moses Mwaniki and Erik Garrison and Nadia Pisanti;
licensed under Creative Commons License CC-BY 4.0
42nd Conference on Very Important Topics (CVIT 2016).
Editors: John Q. Open and Joan R. Access; Article No. XX; pp. XX:1–XX:17
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

where the ith set contains strings of the same length ℓ_i (possibly $= 1$ in the deterministic case) but this length can vary between different sets.
An *indeterminate string* over an alphabet Σ is a sequence of subsets of Σ, and it basically corresponds to the special case of D-string in which $\ell_i = 1$ for all $1 \leq i \leq n$. A great deal of research has been conducted in the bioinformatic literature on this type of uncertain sequence (see [20, 31, 32, 24, 28, 6] and references therein), that are equivalent to a sequence written in the IUPAC notation [21] to represent a position in a DNA sequence that can have multiple possible alternatives. These are commonly used to encode the consensus of a population of sequences [37, 5, 26, 27, 33] in a multiple sequence alignment (MSA).
The more general notion of ED-strings (where within a degenerate position variants can have different sizes), and over them the short read matching problem *elastic-degenerate string matching* (EDSM) problem has attracted some attention in the combinatorial pattern matching community. Since its introduction in 2017 [7], a series of results have been published both for the exact ([38, 14, 25, 15]) as well as for the approximate ([13, 12]) version of the problem.
In this paper we propose a new, fast and exact method to align a string to a D-string, the latter possibly representing an MSA, a pan-genome or a partial assembly. Ours is a base-level heuristic-free alignment method that allows affine gap penalty score functions and the use of arbitrary scores as well as weights. D-strings are one out of many others possible pangenome representation ([40]). Our algorithm exploits the assumption that the similarity the alignment must detect is high enough to suitably adapt the seminal idea of [10] to D-strings, combining at the same time *partial order alignments* ([4]) and the *wavefront* paradigm ([39]) managing to work with affine gap penalty function to score gaps, which is a desirable property when the target of the alignment is to account for INDELs variant events.

This paper is organised as follows: Section 2 gives some preliminary definitions on D-strings and notions of partial order and wavefront alignment, Section 3 describes our algorithm, and Section 4 the experimental validation of the resulting tool dsa we realized.
The implementation of DSA is publicly available at `https://github.com/urbanslug/dsa`, and our D-strings random generator can be found at `https://github.com/urbanslug/simed/`.

2 Preliminary Notions

2.1 D-strings

A *string* X is a sequence of elements on an alphabet Σ, where the *alphabet* Σ is a non-empty finite set of letters of size $|\Sigma|$. The set of all finite strings over an alphabet Σ, including the *empty string* ε of length 0, is denoted by Σ^*, while Σ^+ denotes the set $\Sigma^* \setminus \{\epsilon\}$. The set of all strings of length $k > 0$ over Σ is denoted by Σ^k. For any string X, we denote by $X[i, j]$ the *substring* of X that *starts* at position i and *ends* at position j. In particular, $X[0, j]$ is the *prefix* of X that ends at position j, and $X[i, |X|]$ is the *suffix* of X that starts at position i, where $|X|$ denotes the *length* of X.

▶ **Definition 1.** *A* degenerate string (D-string) $\hat{S} = \hat{S}\{1\}\hat{S}\{2\}\ldots\hat{S}\{n\}$ *of* length n *over an alphabet* Σ *is a finite sequence of* n *degenerate letters* $\hat{S}\{i\}$*'s. Each degenerate letter* $\hat{S}\{i\}$ *has* width $\ell_i > 0$ *and is a finite non-empty set of* $|\hat{S}\{i\}|$ *strings of the same length* ℓ_i *(i.e.* $\hat{S}\{i\}[1], \ldots, \hat{S}\{i\}[|\hat{S}\{i\}|] \in \Sigma^{\ell_i}$*).*

For any $1 \leq i \leq n$, we remark that a degenerate letter $\hat{S}\{i\}$ with $|\hat{S}\{i\}| = \ell_i = 1$ is just a simple letter of Σ. Whenever this happens, we will say that i is a *solid position* of $\hat{S}$. We

now define some parameters that measure the degeneracy of a D-string.

▶ **Definition 2.** *The* total size N *and* total width $w(\hat{T}) = W$ *of a D-string* $\hat{S}$ *are respectively defined as* $N = \sum_{i=1}^{n} |\hat{S}\{i\}| \cdot \ell_i$ *and* $W = \sum_{i=1}^{n} \ell_i$.

▶ **Example 3.** In the D-string $\hat{T}$ = `GCA[AT/CG]C[G/T]GG[TA/AA/AT]TT`, $\ell_4 = \ell_9 = 2$ while all other ℓ_i's are 1, and $|\hat{S}\{4\}| = |\hat{S}\{6\}| = 2$, $|\hat{S}\{9\}| = 3$ while all the other $|\hat{S}\{i\}|$'s are 1; finally, the solid positions of $\hat{T}$ are $1, 2, 3, 5, 7, 8, 10, 11$.

We remark that a D-string of width W actually represents a set of linear strings of length W, corresponding to all the strings that can be read by making any choice at a degenerate positions. Formally, given a D-string $\hat{T}$ of width W, and a string T with $|T| = W$ being any string T in such set, we say that T belongs to $\hat{T}$ ($T \in \hat{T}$).

▶ **Example 4.** The strings `GCACGCTGGAATT` and `GCAATCTGGTATT` are two of the twelve strings that belong to the D-string $\hat{T}$ = `GCA[AT/CG]C[G/T]GG[TA/AA/AT]TT`.

The ith position $\hat{S}\{i\}$ of a D-string $\hat{S}$ ($1 \leq i \leq n$) should not be confused with its ith width $\hat{S}[i]$ ($1 \leq i \leq W$):

▶ **Definition 5.** *For any D-string* $\hat{S}$*, its* i^{th} *width (for* $1 \leq i \leq W$*)* $\hat{S}[i]$ *is the letter, or the set of alternative letters, that can appear in* $T[i]$ *for string* $T \in \hat{T}$*. For any D-string* $\hat{S}$*, we denote by* $\hat{S}[i_1, i_2]$ *the* D-substring *of* $\hat{S}$ *that starts at width* i_1 *and ends at width* i_2 *with* $1 \leq i_1 \leq i_2 \leq W$.

We will also use the notation s_i for both $|\hat{S}\{i\}|$ and $|\hat{S}[i]|$.

▶ **Example 6.** The D-string $\hat{T}$ = `GCA[AT/CG]C[G/T]GG[TA/AA/AT]TT` has length $n = 11$, size $N = 20$, and width $W = 13$. We have that at width, say, 3 there is a solid position `A` = $\hat{S}[3]$, while at width 4 there is `[A/C]` = $\hat{S}[4]$. Furthermore, the D-substring $\hat{S}[3, 7]$ is `A[AT/CG]C[G/T]`. Here below we show positions and widths for the D-string `CA[AT/CG]C[G/T]GG[TA/AA]T` having length 9, width 11, and size 16.

D-string	C	A	[A	T	/	C	G]	C	[G	/	T]	G	G	[T	A	/	A	A]	T
width	1	2	3	4		3	4	5	6		6	7	8	9	10		9	10	11
position	1	2			3			4		5		6	7			8			9

2.2 Partial Order Alignments

The use of Partial Order (graphs) for sequence Alignments (POA) was introduced in [4] with the purpose of improving the iterative step of progressive ([8]) multiple sequence alignments (MSA), that is the step where a new sequence is added to an MSA (*starting MSA*). Traditionally, this was performed by first reducing the starting MSA to a linear profile, and then aligning it to the new sequence in order to obtain a new MSA (*resulting MSA*), with the downside that the reduction of the starting MSA to a linear profile might carry and propagate alignment errors. The idea behind POA is to perform a direct pairwise dynamic programming alignment between the new sequence and a suitable graph representation of the starting MSA (thus replacing the somewhat lossy linear profile representation), with the outcome of a guarantee that, in the resulting MSA, the optimal alignment of the new sequence with respect to each one of the sequences contributing to the starting MSA is maintained. Such graph representation, the *partial order graph*, basically replaces the linear MSA profile with a DAG whose edges represent a partial order between the letters of the MSA to be used as an alternative to the traditional total order of the rows of an alignment.

2.3 Wave Front Alignments with affine gap penalty

In [11], Gene Myers introduced the linear time and space dynamic programming (DP) alignment for similar strings: when aligning two strings of size N whose distance is known to be upper-bounded by D, rather than computing the whole DP table of size N^2, only a stripe of $\mathcal{OD}$ diagonals are computed, as their similarity guarantees their optimal alignment to lay therein. The result is a $\mathcal{O}(\mathcal{DN})$ algorithm replacing the quadratic one.
The *Affine Gap Penalty* function to score the cost of gaps in alignments evaluates a series of k consecutive gaps as $w(k) = o + k \cdot e$ (where o is the cost of opening the gap, and e that of extending it), rather than the sum of k single gap's costs. Using such gap penalty score is almost mandatory in genomic sequences analysis as this forces the optimal alignment to detect and highlight INDELs variations that typically involve several consecutive nucleotides. In order to restore the optimal substructure of the alignment problem when using affine gap penalty score funcion, the DP algorithm requires three alignment tables instead of one: matrices I, D, M that store the score of the best alignments ending - respectively - with an Insertions, a Deletions, or a (Mis)Match.
In [39], the linear time and space optimization was extended to the exact computation of an optimal pairwise alignment of strings using the affine gap penalty score function. Roughly, for the alignment of two strings of size N and M, the three matrices I, D, M of size $N \times M$ are replaced by *wave front* records $I_{d,k}, D_{d,k}, M_{d,k}$ that store, per each score/distance d and diagonal k, the furthest-reaching offset in k that scores d. The dynamic programming recurrence is then operated on increasing values of d and adjacent diagonals k, up to the final distance.

3 The DSA Algorithm

Let us first formally state the problem we solve:

> String Alignment to D-strings (StoDS)
> **Input:** A D-string $\hat{T}$ of width W and size N, a pattern P of length m, and penalty scores a, x, o, e.
> **Output:** An optimal alignment between P and $\hat{T}$ using scores: a for match, x for mismatch, o for gap opening, and e for gap extension.

In the problem statement above, by optimal alignment we mean the alignment between P and a string $T \in \hat{T}$ (see Section 2.1) that minimizes the distance computed by scoring a for a match, x for a mismatch, and gap affine penalty function with o for gap opening and e for gap extension.

In this section we describe our algorithm DSA that optimally solves StoDS. We start with Section 3.1 where we show how we adapted the *Partial Order Alignment* (POA) framework to work with alignment to D-strings. Then, in Section 3.2 we describe our extension of *Wave Front Alignment* (WFA) to D-strings with and show how we merge the two techniques into our algorithm DSA and analyse its complexity.

3.1 Partial Order Alignment with D-strings

In this section, we show how to use partial order alignment to be able to perform base level alignments with D-strings. Without loss of generality, for the sake of simplifying the notation, within this section we will not use the gap affine penalty score (whose implementation with D-strings will be described in the next section), but rather account a cost g to any gap.

Given a D-string $\hat{T}$ of width W and pattern P of length m, we build a $(W+1) \times (m+1)$ DP table M having a row i per each ith width of $\hat{T}$, and a column per each letter of P (plus the usual first row and first column for their empty prefixes), and that will store in $M[i,j]$ the best score of aligning $P[1,j]$ to $\hat{T}[1,i]$ (an example is shown in Figure 1). When a row falls within a non-solid position having s strings of length ℓ, then the row is not associated with a simple letter, but rather with s of them, and this will be the case for ℓ consecutive rows. The entries of these rows will not contain a single score, but rather a tuple of size s: if row i has letters $a_{i,1}, \ldots, a_{i,h}, \ldots, a_{i,s}$, then $M[i,j]$ is a tuple of s values $< M[i,j]^1, \ldots, M[i,j]^h, \ldots, M[i,j]^s >$ where, for each $1 \leq h \leq s$, a match is accounted if and only if $a_{i,h} = P[j]$.
On top of this tuple representation, a partial order is assumed for the N letters of the D-string: within solid positions, the usual order of sequences applies; in a degenerate position, instead, distinct letters in the same tuple are not comparable, and nor is letter $a_{i,h}$ of tuple $a_{i,1}, \ldots, a_{i,h}, \ldots, a_{i,s_i}$ at row i comparable with the other letters of all the ℓ adjacent rows composing the same degenerate position, except for the ℓ other letters $a_{i',h}$ that are also the hth letter in their tuples. Within all comparable letters, the order corresponds to that of the rows of M.

We use this partial order to drive the alignment along the D-string: wherever the traditional dynamic programming alignment algorithm refers to the previous row and/or previous column, here we refer only to entries that - instead - are preceding according to the partial order defined above.

▶ **Example 7.** For the D-string `AC[GC/AT]A` the partial order is defined by the graph shown on the letters at the rows of the DP table of Fogure 1.

We now formalize with dynamic programming recurrence relations the way we apply the ideas sketched above. For any two letters $x, y \in \Sigma$, we define a function $m(x,y)$ as equal to a if $x = y$, and equal to x otherwise. Denoting with s_i the size of the tuple at row i, we compute $M[i,j]$ for all $1 \leq j \leq m$ distinguishing the following cases:

$s_{i-1} = s_i = 1$. In this case no degeneracy is encountered and the traditional dynamic programming alignment framework applies.

$s_{i-1} = 1$ and $s_i = s > 1$. In this case, row i corresponds to the opening of a degenerate position of s_i strings of length ℓ, this row (like the next $\ell - 1$ rows) represents s_i alternative letters $a_{i,1}, \ldots, a_{i,s_i}$ and each entry $M[i,j]$ will contain tuples of size s_i computed as follows for $1 \leq h \leq s_i$:

$$M[i,j]^h = min \begin{cases} M[i-1,j-1] + m(P[j], a_{i,h}) \\ M[i-1,j] + g \\ M[i,j-1]^h + g \end{cases} \tag{1}$$

$s_{i-1} = s_i = s > 1$. In this case, row i is inside the same degenerate position of the previous row[1], this row still represents s alternative letters $a_{i,1}, \ldots, a_{i,s_i}$ and each entry $M[i,j]$ will contain tuples of size s_i computed as follows for $1 \leq h \leq s_i$:

$$M[i,j]^h = min \begin{cases} M[i-1,j-1]^h + m(P[j], a_{i,h}) \\ M[i-1,j]^h + g \\ M[i,j-1]^h + g \end{cases} \tag{2}$$

[1] Should rows $i-1$ and i fall into distinct degenerate positions that have the same size only by chance, this case would be managed as in the last case below.

	-	A	C	G	T	A
-	0	1	2	3	4	5
A	1	0	1	2	3	4
C	2	1	0	1	2	3
G A	(3,3)	(2,2)	(1,1)	(0,1)	(1,2)	(2,2)
C T	(4,4)	(3,3)	(2,2)	(1,2)	(1,1)	(2,3)
A	5	4	3	2	2	1

Figure 1 Dynamic Programming table for the Alignment of the string `ACGTA` against the D-string `AC[GC/AT]A` partially ordered as shown by oriented edges, and using scores: 0 for match, 1 for mismatch, and 2 for gap.

$s_{i-1} = s > 1$ and $s_i = 1$. In this case, row $i-1$ was the last letter of a degenerate position, and row i is a solid position representing a single letter $\hat{T}[i]$. The entry $M[i,j]$ will contain a value computed as follows:

$$M[i,j] = min \begin{cases} M[i-1,j-1]^1 + m(P[j],\hat{T}[i]) \\ \dots \quad \dots \quad \dots \\ M[i-1,j-1]^{s_{i-1}} + m(P[j],\hat{T}[i]) \\ M[i-1,j]^1 + g \\ \dots \quad \dots \\ M[i-1,j]^{s_{i-1}} + g \\ M[i,j-1] + g \end{cases} \tag{3}$$

$s_{i-1} > 1$ and $s_i = s > 1$ but not necessarily $s_{i-1} = s_i$. In this case, row $i-1$ was the last letter of a degenerate position and row i is the first of a new different degenerate position. This row represents s alternative letters $a_{i,1}, \dots, a_{i,s_i}$ and each entry $M[i,j]$ will contain tuples of size s_i computed as follows for $1 \leq h \leq s_i$:

$$M[i,j]^h = min \begin{cases} M[i-1,j-1]^1 + m(P[j],a_{i,h}) \\ \dots \quad \dots \quad \dots \\ M[i-1,j-1]^{s_{i-1}} + m(P[j],a_{i,h}) \\ M[i-1,j]^1 + g \\ \dots \quad \dots \\ M[i-1,j]^{s_{i-1}} + g \\ M[i,j-1]^h + g \end{cases} \tag{4}$$

For example, Figure 1 shows the DP table for the alignment of the D-string `AC[GC/AT]A` against `ACGTA`.

3.2 Wave Front Alignment with D-strings

In this section we show how to perform Wave Front Alignment of Partially Ordered D-Strings with affine gap penalty score in almost linear time using a D-strings customization of Wave Front Alignment.

Table 1 Notation

Symbol	Description
a	score of match
x	score of mismatch
o	score of gap opening
e	score of gap extending
d	current distance (it will start from 0 and reach the solution of the StoDS problem)
P	query pattern: $P = P_0P_1...P_m$ ($m = \Theta(W)$)
$\hat{T}$	D-string $\hat{T} = \hat{T}[0]\hat{T}[1]...\hat{T}[W]$
WF_d	the wavefront WF_d for d
$\widetilde{M}_d, \widetilde{I}_d, \widetilde{D}_d$	components of WF_d that will be defined for each diagonal k in WF_d.
$\widetilde{M}_{d,k}, \widetilde{I}_{d,k}, \widetilde{D}_{d,k}$	contain the values $\text{offset}_{d,k}$ for each diagonal k in WF_d.
lo_d, hi_d	define the width of WF_d
$\text{offset}_{d,k}$	defines for each diagonal k in WF_d the furthest reaching points with distance d.
$\text{abandoned}_{d,k}$	boolean telling whether furthest point has been reached in diagonal k for distance d.

Similarly to [39], we replace the three affine gap penalty dynamic programming matrices I, D, M of size $W \times m$ (W being the width of $\hat{T}$ and m the length of P) with *wave front* records $\widetilde{I}_d$, $\widetilde{D}_d$, $\widetilde{M}_d$ starting with an initial distance $d = 0$. The value of d can thus only increase along the computation, and for each partial score d, the records $\widetilde{I}_d$ (resp. $\widetilde{D}_d$, $\widetilde{M}_d$) store the following values: lo_d, hi_d and a tuple $\text{offset}_{d,k}$ per each diagonal k in the range defined by lo_d and hi_d. Indeed, the values lo_d, hi_d define WF_d: the set of diagonals that allow to reach alignment score d. Per each partial distance d and diagonal k, $\text{offset}_{d,k}$ stores the furthest-reaching offset in diagonal k that scores d. Initially, $lo_d = hi_d = 0$ and the only single starting diagonal is the main diagonal $k = W - m$.

Table 1 summarizes the symbols and data (structure) we are using.

When using WF_d with D-texts, $\widetilde{I}_{d,k}, \widetilde{D}_{d,k}, \widetilde{M}_{d,k}$ are in general tuples of size s (the number of variants). Observe that (d, k, and) $\text{offset}_{d,k}$, identify a specific diagonal and how far you can go down along it keeping score d, and therefore k and $\text{offset}_{d,k}$ actually precisely define an entry in what would have been the matrices D, I and M. Let us name u the row and v the column of the entry defined by $\text{offset}_{d,k}$. Now, if $\hat{T}[u]$ is a single/solid letter, then the offset of $\widetilde{I}_d, \widetilde{D}_d, \widetilde{M}_d$ is a single value, and else it is a tuple of size s_u: a wave front per each letter $\hat{T}[u]$.

Starting with initial distance $d = 0$ and the sole central diagonal, we compute the wavefront for growing values of d up the final result. The value of d is increased, and new offsets are computed, when the borders of the wave front have been reached: no further match is possible, and either a mismatch score x must be accounted (and new offset computed along the same diagonal k for distance $d + x$), or a gap must be accounted and therefore new diagonals must be explored (and new offset computed for distance that has increased according to the gap penalty function).

Let us now formalize the Wave Front Alignment for D-strings with the recurrence relations that show how to update, for each diagonal k, the offsets information, We will use the generic notation $\text{offset}_{d,k}$ as well as the more specific notation $\text{offset}^D_{d,k}$ (resp. $\text{offset}^I_{d,k}$, $\text{offset}^M_{d,k}$) to

mean the offsets in $\widetilde{D}_d$ (resp. $\widetilde{I}_d$, $\widetilde{M}_d$).
The formulae below are basically those of [39] with the tuple information added. When computing $\widetilde{I}_{d,k}$ we are assuming that the last event has been an insertion, in which case the previous event could either have been (i) a (mis)match, in which case a new gap is being opened, or (ii) another insertion, in which case an existing gap is being extended. In case (i) the new offset is that of $\widetilde{M}_{d-o-e}$ because the distance is increased by $o+e$, while in case (ii) the new offset is that of $\widetilde{I}_{d-e}$ because the distance is only increased by the gap extension e. In both cases, the new offset has to be picked from the preceding diagonal $k-1$. The following formula formalizes this, describing how to compute the new offsets for $\widetilde{I}_{d,k}$, assuming the general case in which the entry on which we do the recurrence contains a tuple of size s:

$$offset^I_{d,k} = 1 + max \begin{cases} \max\limits_{\texttt{in tuple of size s}} (offset^M_{d-o-e,k-1}) & \text{(Open insertion)} \\ \max\limits_{\texttt{in tuple of size s}} (offset^I_{d-e,k-1}) & \text{(Extend insertion)} \end{cases} \tag{5}$$

Dually, $\widetilde{D}_{d,k}$ is computed assuming that the last event has been a deletion, in which case the previous event could either have been a (mis)match, in which case a new gap is being opened, or another deletion, in which case an existing gap is being extended. The new offset is then either in $\widetilde{M}_{d-o-e}$ or in $\widetilde{D}_{d-e}$, and in both cases the new offset has to be picked from the next diagonal $k-1$. The following formula formalizes this, describing how to compute the new offsets for $\widetilde{D}_{d,k}$, assuming the general case in which the entry on which we do the recurrence contains a tuple of size s:

$$offset^D_{d,k} = max \begin{cases} \max\limits_{\texttt{in tuple of size s}} (offset^M_{d-o-e,k+1}) & \text{(Open deletion)} \\ \max\limits_{\texttt{in tuple of size s}} (offset^D_{d-e,k+1}) & \text{(Extend deletion)} \end{cases} \tag{6}$$

Dually, $\widetilde{M}_{d,k}$ is computed assuming that the last event has been a (mis)match, in which case we stay in the same diagonal k, and the previous event could either have been another (mis)match, or an insertion or deletion. The new offset is then either in $\widetilde{M}_{d-x}$ or in $\widetilde{D}_d$ or $\widetilde{I}_d$. The following formula formalizes this, describing how to compute the new offsets for $\widetilde{M}_{d,k}$, assuming the general case in which the entry on which we do the recurrence contains a tuple of size s:

$$offset^M_{d,k} = max \begin{cases} offset^I_{d,k} & \text{(Insertion)} \\ offset^D_{d,k} & \text{(Deletion)} \\ \max\limits_{\texttt{in tuple of size s}} (offset^M_{d-x,k}) + 1 & \text{(Mismatch)} \end{cases} \tag{7}$$

In the formula for $\widetilde{M}_{d,k}$, the case of a match is not taken into account because offsets are recomputed only when a match is not available.

Notice that, among possible values, the maximum is always sought here because the distance is a fixed parameter d in $\widetilde{M}_{d,k}, \widetilde{I}_{d,k}, \widetilde{D}_{d,k}$, and with that one wants to go as far as possible along k maintaining that distance (that can only increase otherwise), therefore eventually minimizing the final d whose WF_d will reach the bottom right final entry of the matrix.

In all formulae above, whenever incurring in a gap and increasing or decreasing the diagonal number, then this has to be intended as updating lo_d or hi_d.

The pseudocode of DWF_ALIGN is shown in Algorithm 1 (structured as in [39] but modified to be adapted to D-strings and their tuples based partial order alignment realization). The function DWF_EXTEND (Algorithm 2) is called to extend the wavefront extends $\widetilde{M}_{d,k}$ for all diagonals k between lo_d and hi_d, and it calls the function λ (Algorithm 3), which is the core of DSA . This function is where the offsets are actually increased. When this happens, then λ:
(i) Updates the width of the offset tuple with that of the current row to account for the case in which we extend a wavefront from an index i by 1 into an index $i+1$ such that $|\widehat{T}[i]| \neq |\widehat{T}[i+1]|$ (the size of the tuple changes), that happens when the wavefront has been extended into a degenerate letter or out of a degenerate letter. In this case, we need to increase or reduce the length of offsets to correspond to the number of variants at $\widehat{T}[i+1]$. We do this by selecting $max(\widehat{T}[i])$ and setting it as the offset for each variant at $\widehat{T}[i+1]$.
(ii) Picking the maximum from previous row does not incur into an artefact match[2] thanks to the boolean vector *abandoned* (if that part of a variant cannot be extended any further, then we mark as abandoned the extension of the wavefront for the given score) that will allow to know along which variant we can extend the wavefront, and will also be used by the traceback function to compute the correct alignment

Algorithm 1

```
function DWF_ALIGN(T̂, P, λ)
    W ← |w(T̂)|                                   // width of the D-Text
    A_k ← |T̂| − |P|                               // central diagonal
    end ← max{W, |P|}                             // end of alignment index
    d ← 0; s_0 ← |T̂[0]|                           // initial edit distance and width
    offset^M_{0,A_k} ←< 0, ..., 0 >               // tuple of length s_0 initialized to 0s
    abandoned^M_{0,A_k} ←< false, ..., false >    // tuple of length s_0 initialized to false
    while true do
        // extend the M-wavefront
        DWF_EXTEND(M̃, λ, d)
        f ← max in tuple offsets^M_{d,A_k}
        if f ≥ end then break;                    //  reached end of alignment
        d ← d + 1                                 // increase current distance
        DWF_NEXT(M̃, Ĩ, D̃, d)                      // compute the bounds of the next wavefront
    A ← DWF_TRACEBACK(M̃, Ĩ, D̃, d)
    // return the edit distance and the alignment
    return (d, A)
```

The final result we have obtained is an exact algorithm that computes the edit distance (or any distance accounting for insertions, deletions, and substitutions with any desired score, including affine gap penalty score), between a string and a D-string in time and space proportional to $D \cdot N$, where N id the size of the D-string, and D is the computed distance.

We remark that both the optimization problem we addressed and the dynamic programming solution we suggested can naturally make use of weights to be added to input data to be included in the distance function to be optimised. These could be confidence value for

[2] With artefact match we mean to, say, align exactly pattern $P = ACGT$ to $\widehat{T} = A\{TC, GA\}T$ matching the two letters CG that belong to distinct variants.

Algorithm 2

```
function DWF_EXTEND(M̃, λ, d)
    for k from M̃_d^lo to M̃_d^hi do
        ab ← abandoned_{d,k}^M                    // tuple of abandoned variants
        of ← offsets_{d,k}^M                      // tuple of offsets
        f ← max in tuple of                       // index in T̂
        u ← f                                     // index in P
        v ← f − k                                 // extend WF_d
        while λ(v, u, of, ab) do
```

Algorithm 3

```
function λ(v, u, offsets, abandoned)
    if u < 0 or v < 0 or v ≥ |P| or u ≥ W then return false
    s_prev ← length(offsets)                      // size of previous tuple
    s_curr ← |T̂[u]|                               // size of current tuple
    // we are stepping into or out of a degenerate letter
    if s_curr ≠ s_prev then
        // furthest offset
        f ← max in tuple offsets
        // create tuple of size s_curr with furthest offset
        offsets ←< f, ..., f >
        // create tuple of s_curr size with false
        abandoned ←< false, ..., false >
    extend ← false
    previously_extended ← false
    // attempt to extend each variant by 1
    for s′ from 0 to s_curr do
        if T̂[h][s′] = P[v] and abandoned[s′] = false then
            offsets[s′] ← offsets[s′] + 1
            extend ← true
            if previously_extended = false then
                v ← v + 1
                u ← h + 1
                previously_extended ← true
        if T̂[u][s′] ≠ P[v] then
            abandoned[s′] ← true
    return extend
```

bases as well as multiplicity values in the D-string (accounting for variants that are more frequent than others).

We implemented our algorithm in a prototype tool dsa that computes the edit distance and an optimal alignment (thus keeping all the information needed for the tracing back the optimal path). In the next section we show some first comparative experimental validation of the performances of DSA .

4 Experimental validation

4.1 State of the art

To the best of our knowledge, there is no published software tool implementing an algorithm specifically designed for a global alignment of a linear string and a D-text. Some algorithm are designed for *semi*global alignment, that is with the pattern being (substantially) shorter than the D-string (e.g. [9, 22, 3, 1, 2]); some of them and others are designed for more general purposes, such as aligning a linear string to a graph structure that generalizes D-strings (e.g.[9, 22, 2, 30, 29, 19]). Some of these and other tools are not exact and hence do not guarantee to find the optimal alignment or solution (e.g [29, 30, 19, 16, 17, 18]). Finally, some algorithms that basically solve a very similar problem, actually do not compute any distance or alignment, but rather output a consensus string that suitably merges the input strings (e.g. [41, 34, 36, 35]).

The Variation Graph Toolkit VG suite [9, 22] contains a tool that can solve StoDS aligning a string to a more general data structure than D-strings (Variation Graphs indeed), but this is specifically designed to map reads that are (much) shorter that the graph/text, rather than making a global alignment, and therefore a comparison would not be fair. The same holds for SOPANG and its extension SOPANG2 ([3, 1]): they are designed for *short* reads mapping on ED-strings, and moreover they only detect exact matches without allowing gaps nor mismatches. Also the tool VARGAS ([2]) works with short reads.
We thus considered the tool abPOA ([41]), a C library tool to align a sequence to a directed acyclic graph that also uses partial order alignment and supports global alignment, but we could not run in on data of length $100,000$b (accuracy comparison would not have been possible anyhow, as abPOA does not output an alignment nor a distance, but rather a consensus string. The same (could not run on length $100,000$b strings and output is a consensus string) holds for the tool Astarix ([34, 36, 35]), which addresses the string to graph alignment by solving a shortest path problem on a suitably designed graph representation of input data.

We therefore compared the performances of DSA with those of GraphAligner ([29, 30]) and Minigraph ([19, 16, 17, 18]) for solving StoDS. Both of them are designed to align strings on a more general graph structure that D-strings and therefore the comparison we show below should be intended as a validation of the performances of DSA in solving StoDS, and not as DSA being in general a better tool than any of the other two.
GraphAligner is designed to align reads to genome graphs as a key to many applications, including error correction, genome assembly, and genotyping of variants in a pangenome graph. In aligning long reads to genome graphs, GraphAligner was 13x faster and used 3x less memory when compared in [29] to the state-of-the-art tools. As for accuracy, when employing GraphAligner for error correction, in [29] it turned out more than twice as accurate and over 12x faster than extant tools. GraphAligner uses a seed and extend

method and the bitvector alignment extension algorithm of [30].
MINIGRAPH is a sequence-to-graph mapper and graph constructor. It aligns a query sequence against a sequence graph using a gap affine penalty score. MINIGRAPH is a well maintained and highly optimized software tool that uses minimizers to find strong colinear chains as starting point to build the alignment ([19, 16, 17, 18]).

4.2 Experiments setup

We randomly generated a D-string of width $W = 100,000$b by first generating a random string of length W on alphabet $\{A, C, G, T\}$, and then inserting[3] therein g (input parameter to our random generator of D-strings given as a percentage of W) degenerate non-solid positions as follows: for each such position p we pick at random a value for its size s_p between 1 and S (another input parameter), and a randomly chosen length ℓ_p between 1 and L (input parameter again).
We generated D-strings using width $W = 100,000$b in all tests, degeneracy frequencies values $g = 1\%, 10\%, 20\%$, maximum variance values $S = 2, 5$, and maximum variant lengths $L = 1, 4$. As a consequence the width of tested D-strings will always be $W = 100,000$b, while its total size N will be greater than W and depending from input parameters g, S, L.
To this purpose, we implemented the random D-strings generator SIMED, which is publicly available at `https://github.com/urbanslug/simed/`.
From the obtained synthetic D-string $\hat{T}$, we extracted a ground truth exact pattern P_0 of size W (that is, a string $P \in \hat{T}$ that thus matches $\hat{T}$ with distance 0), and we (possibly) modified P_0 into the actual input query P with different possible divergences using real .vcf files.
The divergences we tested were to insert no divergence at all, or $0, 1\%$ or 1% SNPs, or $0, 1\%$ INDELs (percentages are on W). Hence, the size $m = |P|$ of the query string will be in $\Theta(W)$ and so will the distance d between P and $\hat{T}$.

4.3 Preliminary Results

All tests were ran on a laptop (single threaded) Intel® Core™ i7-11800H × 16 with 16.0 GiB RAM. Space and time was measured using `/usr/bin/time -f"%S\t%M"` to extract system time (seconds) and maximum resident set size (kbytes). In all tests below we used alignment scores $a = 0, x = 1, o = 2, e = 1$.

In all experiments, for all tools time was below 0.1 seconds and therefore negligible and not reported case per case. We do report, instead the memory peak (weakness of the current prototype of DSA : as we mention in further work session, we plan to improve that), the number of mistakes in the alignment found (only for the cases with no INDELs because there the alignment is affected by differences in penalty scores), and finally the last column shows the number of detected events on optimal alignment: number M of matches, X for mismatches, I for insertions, and D for deletions. When the pattern divergence is only SNPs, then for a correct alignment it must be $I = D = 0$, X approximately equal to the number of SNPs (except where by chance the divergence will not change the DNA base), and $M = W - X$. For the experiments involving INDELs divergence in the pattern, we also report the number G of gaps that are opened: with $W = 100,000$b and $0.1\% INDELs$, in a correct alignment it must be $G = 100$.

[3] The insertion was done forcing the width to remain W

Table 2 below shows results with a D-string of size $N = 101,000$ generated with lowly frequent low degeneracy (only $g = 1\%$ of positions are degenerate, with at most $S = 2$ variants of length $L = 1$), and with four different pattern divergence: no divergence, 0.1% of SNPs, 1% of SNPs, and 0.1% of INDELs.

Table 2 Performances with D-string of size $N = 101,000$ with degeneracy $g = 1\%$, $S = 2$, $L = 1$.

Tool	Pattern divergence	Peak memory (kbytes)	Total Error (absolute)	Alignment events
DSA	none	3162	0	100,000M 0X 0I 0D
GraphAligner	"	6244	0	100,000M 0X 0I 0D
Minigraph	"	5607	9	99,991M 9X 0I 0D
DSA	0.1 % SNPs	42209	0	99902M 98X 0I 0D
GraphAligner	"	23972	1	99903M 97X 0I 0D
Minigraph	"	8058	2	99891M 100X 0I 0D
DSA	1 % SNPs	474963	0	99003M 997X 0I 0D
GraphAligner	"	23960	10	99001M 995X 4I 4D
Minigraph	"	7928	3	98991M 1000X 0I 0D
DSA	0.1 % INDELs	99322	0	99865M 0X 106I 133D 100G
GraphAligner	"	23973	-	99757M 111X 105I 132D 140G
Minigraph	"	8162	-	99865M 0X 106I 133D 100G

Table 3 below shows results with a D-string of size $N = 106,147$ generated with lowly frequent high degeneracy (only $g = 1\%$ of positions are degenerate, with up to $S = 5$ variants of length up to $L = 4$), and with four different pattern divergence: no divergence, 0.1% of SNPs, 1% of SNPs, and 0.1% of INDELs.

Table 3 Performances with D-string of size $N = 106,147$ with degeneracy $g = 1\%, S = 5, L = 4$.

Tool	Pattern divergence	Peak memory (kbytes)	Total Error (absolute)	Alignment events
DSA	none	3188	0	100,000= 0X 0I 0D
GraphAligner	"	7418	0	100,000= 0X 0I 0D
Minigraph	"	5585	11	99989= 0X 0I 0D
DSA	0.1 % SNPs	44096	0	99900= 100X 0I 0D
GraphAligner	"	26799	1	99901M 99X 0I 0D
Minigraph	"	8038	11	99889M 100X 0I 0D
DSA	1 % SNPs	475668	0	99007= 993X 0I 0D
GraphAligner	"	26682	6	99005M 993X 2I 2D
Minigraph	"	7910	11	98989M 1000X 0I 0D
DSA	0.1 % INDELs	131311	0	99813M 0X 135I 174D 100G
GraphAligner	"	26816	-	99717M 109X 135I 174D 145G
Minigraph	"	8213	-	99813M 1X 136I 175D 100G

Table 4 below shows results with a D-string of size $N = 110,000$ generated with medium frequent low degeneracy ($g = 10\%$ of positions are degenerate, with only up to $S = 2$ variants of length $L = 1$), and with four different pattern divergence: no divergence, 0.1% of SNPs,

1% of SNPs, and 0.1% of INDELs.

Table 4 Performances with D-string of size $N = 110,000$ with degeneracy $g = 10\%, S = 2, L = 1$.

Tool	Pattern divergence	Peak memory (kbytes)	Total Error (absolute)	Alignment events
DSA	none	3174	0	100,000M 0X 0I 0D
GraphAligner	"	16405	0	100,000M 0X 0I 0D
Minigraph	"	5609	6	99994M 0X 0I 0D
DSA	0.1 % SNPs	46942	0	99903M 97X 0I 0D
GraphAligner	"	50386	0	99903M 97X 0I 0D
Minigraph	"	8078	12	99894M 100X 0I 0D
DSA	1 % SNPs	456316	0	99022M 978X 0I 0D
GraphAligner	"	50250	12	99020M 976X 4I 4D
Minigraph	"	7934	94	98947M 997X 0I 0D
DSA	0.1 % (INDELS)	105904	0	99806M 0X 171I 183D 100G
GraphAligner	"	50223	-	99708M 109X 171I 183D 160G
Minigraph	"	8145	-	99806M 0X 171I 183D 100G

Table 5 below shows results with a D-string of size $N = 160,327$ generated with medium frequent high degeneracy ($g = 10\%$ of positions are degenerate, with up to $S = 5$ variants of length up to $L = 4$), and with four different pattern divergence: no divergence, 0.1% of SNPs, 1% of SNPs, and 0.1% of INDELs.

Table 5 Performances with D-string of size $N = 160,327$ with degeneracy $g = 10\%, S = 5, L = 4$.

Tool	Pattern divergence	Peak memory (kbytes)	Total Error (absolute)	Alignment events
DSA	none	3134	0	100,000= 0X 0I 0D
GraphAligner	"	27417	0	100,000= 0X 0I 0D
Minigraph	"	5536	4	99996= 0X 0I 0D
DSA	0.1 % SNPs	66656	0	99905= 95X 0I 0D
GraphAligner	"	76167	0	99905M 95X 0I 0D
Minigraph	"	8046	14	99896M 100X 0I 0D
DSA	1 % SNPs	469600	0	99022= 978X 0I 0D
GraphAligner	"	76129	13	99024M 973X 3I 3D
Minigraph	"	7878	48	98804M 1000X 0I 0D
DSA	0.1 % (INDELS)	126051	0	99814M 0X 108I 186D 100G
GraphAligner	"	76626	-	99665M 147X 110I 188D 170G
Minigraph	"	8199	-	99804M 0X 108I 186D 100G

Table 6 below shows results with a D-string of size $N = 120,000$ generated with frequent but low degeneracy ($g = 20\%$ of positions are degenerate, with only up to $S = 2$ variants of length $L = 1$), and with four different pattern divergence: no divergence, 0.1% of SNPs, 1% of SNPs, and 0.1% of INDELs.

In all tables, the first experiment is for an exact match: no divergence is introduced in the pattern, and an exact match is supposed to be found. As we can see above, DSA always does find the exact solution with 100000 matches, 0 mismatches, and 0 gaps (like GraphAligner

Table 6 Performances with D-string of size $N = 120,000$ with degeneracy $g = 20\%, S = 2$ and $L = 1$.

Tool	Pattern divergence	Peak memory (kbytes)	Total Error (absolute)	Alignment events
DSA	none	3166	0	100000M 0X 0I 0D
GraphAligner	"	21224	0	100000M 0X 0I 0D
Minigraph	"	5604	6	99994M 0X 0I 0D
DSA	0.1 % SNPs	53252	0	99901M 99X 0I 0D
GraphAligner	"	76291	2	99902M 98X 0I 0D
Minigraph	"	8078	8	99894M 100X 0I 0D
DSA	1 % SNPs	436404	0	99048M 952X 0I 0D
GraphAligner	"	76002	6	99047M 951X 2I 2D
Minigraph	"	7913	102	98994M 1000X 0I 0D
DSA	0.1 % (INDELS)	103666	0	99908M 0X 96I 88D 100G
GraphAligner	"	76305	-	99813M 91X 150I 96D 141G
Minigraph	"	8132	-	99908M 0X 96I 88D 100G

does) taking less memory than Minigraph and much less than GraphAligner.
In the second and third experiments we introduced SNPs (in 0.1% and 1% of the positions, respectively): in all experiments, DSA is the only one which is exact (except in two cases where for 0.1% SNPs also GraphAligner is correct but with higher memory usage) at a cost of the highest memory consumption with 1% SNPs, and only higher than Minigraph's for 0.1% SNPs (but therein Minigraph is always the least accurate).
Finally, with INDELs, we can see that the memory usage of DSA must be improved, but our prototype is always exact (and Minigraph almost is, as only in one case it detects a spurious mismatch). GraphAligner is almost always wrong which is expected given it is more general and does not aim to find an exact alignment.

Conclusions and further work

As mentioned above, at the moment our implementation of DSA is just a promising prototype. We plan to improve its memory consumption using the ideas of [23]. Also, we are working on improve its speed (as well as consequent memory use) by removing hopeless diagonals suffixes that are currently being kept in the range.
Also, we remark that with our dynamic programming method, it is very natural to add weights to letters and to use a sum of weights modification of the actual score as objective function, in order to account for possible useful metadata such as confidence level in the query string or D-strings bases, or abundance in the MSA represented by the D-string.

References

1 A.Cisłak and Szymon S.Grabowski. SOPanG 2: online searching over a pan-genome without false positives. *arXiv:2004.03033 [cs]*, 2020. URL: `https://arxiv.org/abs/2004.03033`.

2 C.A.Darby, R.Gaddipati, M.C.Schatz, and B.Langmead. Vargas: heuristic-free alignment for assessing linear and graph read aligners. *Bioinformatics*, 36(12):3712–3718, 2020.

3 Aleksander Cisłak, Szymon Grabowski, and Jan Holub. SOPanG: online text searching over a pan-genome. *Bioinformatics*, 34(24):4290–4292, 06 2018. `arXiv:https://`

academic.oup.com/bioinformatics/article-pdf/34/24/4290/27088655/bty506.pdf, doi:10.1093/bioinformatics/bty506.

4 C.Lee, C.Grasso, and M.F.Sharlow. Multiple Sequence Alignment Using Partial Order Graphs. *Bioinformatics*, 18(3):452–464, 2002.

5 The Computational Pan-Genomics Consortium. Computational Pan-Genomics: Status, Promises and Challenges. *Briefings in Bioinformatics*, 19(1):118–135, 2018.

6 C.S.Iliopoulos and J.Radoszewski. Truly Subquadratic-Time Extension Queries and Periodicity Detection in Strings with Uncertainties. In *27th Annual Symposium on Combinatorial Pattern Matching (CPM)*, volume 54 of *LIPIcs*, pages 8:1–8:12, 2016.

7 C.S.Iliopoulos, R.Kundu, and S.P.Pissis. Efficient Pattern Matching in Elastic-Degenerate Texts. In *11th International Conference on Language and Automata Theory and Applications (LATA)*, volume 10168 of *Springer LNCS*, pages 131–142, 2017.

8 D.-F.Feng and R.F. Doolittle. Progressive sequence alignment as a prerequisite to correct phylogenetic trees. *Journal of Molecular Evolution*, 25(4):351–360, 1987.

9 E.Garrison, J.Sirén, A.M.Novak, G.Hickey, J.M.Eizenga, E.T.Dawson, W.Jones, S.Garg, C.Markello, M.F.Lin, B.Paten, and R.Durbin. Variation Graph Toolkit Improves Read Mapping by Representing Genetic Variation in the Reference. *Nature Biotechnology*, 36(9):875–879, 2018.

10 E.W.Myers. AnO(ND) difference algorithm and its variations. *Algorithmica*, 1(1):251–266, November 1986. Number: 1. doi:10.1007/BF01840446.

11 E.W.Myers. An O(ND) difference algorithm and its variations. *Algorithmica*, 1(2):251–266, 1986.

12 G.Bernardini, N.Pisanti, S.P.Pissis, and G.Rosone. Pattern matching on elastic-degenerate text with errors. In *24th International Symposium on String Processing and Information Retrieval (SPIRE)*, volume 10508 of *Springer LNCS*, pages 74–90, 2017.

13 G.Bernardini, N.Pisanti, S.P.Pissis, and G.Rosone. Approximate pattern matching on elastic-degenerate text. *Theoretical Computer Science*, 812:109–122, April 2020.

14 G.Bernardini, P.Gawrychowski, N.Pisanti, S.P.Pissis, and G.Rosone. Even Faster Elastic-Degenerate String Matching via Fast Matrix Multiplication. In *46th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 132 of *LIPIcs*, pages 21:1–21:15, 2019.

15 G.Bernardini, P.Gawrychowski, N.Pisanti, S.P.Pissis, and G.Rosone. Elastic-Degenerate String Matching via Fast Matrix Multiplication. *SIAM Journal of Computing*, 51(3):549–576, 2022.

16 H.Li. Minimap and Miniasm: Fast Mapping and de Novo Assembly for Noisy Long Sequences. *Bioinformatics*, 32(14):2103–2110, 2016.

17 H.Li. Minimap2: Pairwise Alignment for Nucleotide Sequences. *Bioinformatics*, 34(18):3094–3100, 2018.

18 H.Li. New strategies to improve minimap2 alignment accuracy. *Bioinformatics*, 37(23):4572–4574, 2021.

19 H.Li, X.Feng, and C.Chu. The Design and Construction of Reference Pangenome Graphs with Minigraph. *Genome Biology*, 21(265), 2020.

20 H.Soldano, A.Viari, and M.Champesme. Searching for flexible repeated patterns using a non-transitive similarity relation. *Pattern Recognition Letters*, 16(3):233–246, 1995.

21 IUPAC-IUB Commission on Biochemical Nomenclature. Abbreviations and symbols for nucleic acids, polynucleotides, and their constituents. *Biochemistry*, 9(20):4022–4027, 1970.

22 J.M.Eizenga, A.M.Novak, E.Kobayashi, F.Villani, C.Cisar, S.Heumos, G.Hickey, V.Colonna, B.Paten, and E.Garrison. Efficient dynamic variation graphs. *Bioinformatics*, 36(21):5139–5144, 2021.

23 J.M.Eizenga and B.Paten. Improving the time and space complexity of the wfa algorithm and generalizing its scoring. *bioRxiv*, 2022. URL: https://www.biorxiv.org/content/early/2022/01/15/2022.01.12.476087.

24 K.Abrahamson. Generalized string matching. *SIAM Journal of Computing*, 16(6):1039–1051, 1987.

25 K.Aoyama, Y.Nakashima, T.I, S.Inenaga, H.Bannai, and M.Takeda. Faster Online Elastic Degenerate String Matching. In *29th Annual Symposium on Combinatorial Pattern Matching (CPM)*, volume 105 of *LIPIcs*, pages 9:1–9:10, 2018.

26 M.Alzamel, L.A.K.Ayad, G.Bernardini, R.Grossi, C.S.Iliopoulos, N.Pisanti, S.P.Pissis, and G.Rosone. Degenerate String Comparison and Applications. In *18th International Workshop on Algorithms in Bioinformatics (WABI)*, volume 113 of *LIPIcs*, pages 21:1–21:14, 2018.

27 M.Alzamel, L.A.K.Ayad, G.Bernardini, R.Grossi, C.S.Iliopoulos, N.Pisanti, S.P.Pissis, and G.Rosone. Comparing Degenerate Strings. *Fundamenta Informaticae*, 175(1-4):41–58, 2020.

28 M.Crochemore, C.S.Iliopoulos, T.Kociumaka, J.Radoszewski, W.Rytter, and T.Walen. Covering problems for partial words and for indeterminate strings. *Theoretical Computer Science*, 698:25–39, 2017.

29 M.Rautiainen and T.Marschall. Graphaligner: rapid and versatile sequence-to-graph alignment. *Genome Biology*, 21(253), 2020.

30 M.Rautiainen, V.Mäkinen, and T.Marschall. Bit-parallel sequence-to-graph alignment. *Bioinformatics*, 35(19):3599–3607, 2019.

31 N.Pisanti, H.Soldano, and M.Carpentier. Incremental inference of relational motifs with a degenerate alphabet. In *16th Annual Symposium on Combinatorial Pattern Matching (CPM)*, volume 3537 of *Springer LNCS*, pages 229–240, 2005.

32 N.Pisanti, H.Soldano, M.Carpentier, and J.Pothier. A relational extension of the notion of motifs: Application to the common 3d protein substructures searching problem. *Journal of Computational Biology*, 16(12):1635–1660, 2009.

33 P.Gawrychowski, S.Ghazawi, and G.M.Landau. On indeterminate strings matching. In *31st Annual Symposium on Combinatorial Pattern Matching (CPM)*, volume 161 of *LIPIcs*, pages 14:1–14:14, 2020.

34 P.Ivanov, B.Bichsel, H.Mustafa, A.Kahles, G.Rätsch, and M.T.Vechev. Astarix: Fast and optimal sequence-to-graph alignment. In *24th Annual International Conference on Research in Computational Molecular Biology (RECOMB)*, volume 12074 of *Springer LNCS*, pages 104–119, 2020.

35 P.Ivanov, B.Bichsel, and M.T.Vechev. Fast and optimal sequence-to-graph alignment guided by seeds. In *26th Annual International Conference on Research in Computational Molecular Biology (RECOMB)*, Springer LNCS In Press, 2022.

36 P.Ivanov, B.Bichsel, and M.Vechev. Fast and optimal sequence-to-graph alignment guided by seeds. *bioRxiv*, 2022. URL: `https://www.biorxiv.org/content/10.1101/2021.11.05.467453v2`.

37 P.Peterlongo, N.Pisanti, F.Boyer, A.Pereira do Lago, and M.-F.Sagot. Lossless filter for multiple repetitions with hamming distance. *Journal of Discrete Algorithms*, 6(3):497–509, 2008.

38 R.Grossi, C.S.Iliopoulos, C.Liu, N.Pisanti, S.P.Pissis, A.Retha, G.Rosone, F.Vayani, and L.Versari. On-Line Pattern Matching on Similar Texts. In *28th Annual Symposium on Combinatorial Pattern Matching (CPM)*, volume 78 of *LIPIcs*, pages 9:1–9:14, 2017.

39 S.Marco-Sola, J.C.Moure, M.Moreto, and A.Espinosa. Fast Gap-Affine Pairwise Alignment Using the Wavefront Algorithm. *Bioinformatics*, 37(4):456–463, 2021.

40 V.Carletti, P.Foggia, E.Garrison, L.Greco, P.Ritrovato, and M.Vento. Graph-based representations for supporting genome data analysis and visualization: Opportunities and challenges. In *12th International Workshop on Graph-Based Representations in Pattern Recognition (GbRPR)*, volume 11510 of *Springer LNCS*, pages 237–246, 2019.

41 Y.Gao, Y.Liu, Y.Ma, B.Liu, Y.Wang, and Y.Xing. abPOA: an SIMD-based C library for fast partial order alignment using adaptive band. *Bioinformatics*, 37(15):2209–2211, 2021.