

CONWAY-BROMAGE-LYNDON (CBL): AN EXACT, DYNAMIC REPRESENTATION OF K-MER SETS

Igor MARTAYAN, Bastien CAZAUX, Antoine LIMASSET & Camille MARCHET

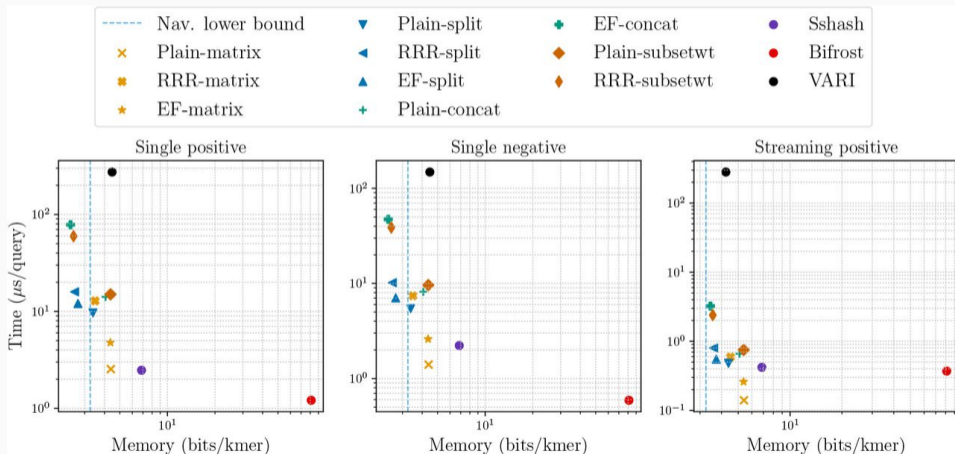
March 15, 2024

DSB 2024 — Montpellier



MOTIVATION OF THIS WORK

Plenty of compact data structures for storing k -mers ...but most of them are **static**



Query time and memory usage of some efficient data structures, taken from [Alanko et al. 22]

Goal: designing a **dynamic** index of k -mers
with fast queries and relatively good compression

- membership
- enumeration
- **insertion**
- **deletion**
- set operations (\cup, \cap, \setminus)

CTGAAATG...

CTGAA

TGAAA

GAAAT

AAATG

batch queries



STARTING FROM A SIMPLE IDEA: k -MERS AS A SPARSE SET OF INTEGERS

- we can see k -mers as integers in $\llbracket 4^k \rrbracket$
 $A \rightarrow 00$ $C \rightarrow 01$ $G \rightarrow 10$ $T \rightarrow 11$
- since they're usually very sparse,
we can store them in a sparse bitvector
(as in [Conway & Bromage 11])

Limitations:

- difficult to compress
(especially if it's dynamic)
- not cache-efficient
 $\text{id}(\text{ATGGCA}) \ll \text{id}(\text{TGGCAT})$
(average distance of $4^k/3$)

STARTING FROM A SIMPLE IDEA: k -MERS AS A SPARSE SET OF INTEGERS

- we can see k -mers as integers in $\llbracket 4^k \rrbracket$
 $A \rightarrow 00$ $C \rightarrow 01$ $G \rightarrow 10$ $T \rightarrow 11$
- since they're usually very sparse,
we can store them in a sparse bitvector
(as in [Conway & Bromage 11])

Limitations:

- difficult to compress
(especially if it's dynamic)
- not cache-efficient
 $\text{id}(\text{ATGGCA}) \ll \text{id}(\text{TGGCAT})$
(average distance of $4^k/3$)

What if we changed our representation of k -mers?

THE NECKLACE TRANSFORMATION

NECKLACE TRANSFORMATION OF K-MERS



The **necklace** of x is its **smallest cyclic rotation** $\langle x \rangle = \min_{0 \leq i < k} x^{(i)}$

To make this transformation reversible, keep track of the rotation index

$$x \mapsto (\langle x \rangle, \text{rotation index})$$

NECKLACES OF CONSECUTIVE k -MERS SHARE LONG PREFIXES

k -mer view

GTCGTTCTTCCTAACGTCATCTCTCATTCTG
TCGTTCTTCCTAACGTCATCTCTCATTCTGT
CGTTCTTCCTAACGTCATCTCTCATTCTGTG
GTTCTTCCTAACGTCATCTCTCATTCTGTGA
TTCTTCCTAACGTCATCTCTCATTCTGTGAC
TCTTCCTAACGTCATCTCTCATTCTGTGACA
CTTCCTAACGTCATCTCTCATTCTGTGACAC
TTCCTAACGTCATCTCTCATTCTGTGACAG
TCCTAACGTCATCTCTCATTCTGTGACACGC
CCTAACGTCATCTCTCATTCTGTGACACGCA
CTAACGTCATCTCTCATTCTGTGACACGCAG
TAACGTCATCTCTCATTCTGTGACACGCAGG
AACGTCATCTCTCATTCTGTGACACGCAGGG
ACGTCATCTCTCATTCTGTGACACGCAGGGT

NECKLACES OF CONSECUTIVE k -MERS SHARE LONG PREFIXES

necklace view

AACGTCATCTCTCATTCTG GTCGTTCTTCCT
AACGTCATCTCTCATTCTGT TCGTTCTTCCT
AACGTCATCTCTCATTCTGTG CGTTCTTCCT
AACGTCATCTCTCATTCTGTGA GTTCTTCCT
AACGTCATCTCTCATTCTGTGAC TTCTTCCT
AACGTCATCTCTCATTCTGTGACA TCTTCCT
AACGTCATCTCTCATTCTGTGACAC CTTTCCT
AACGTCATCTCTCATTCTGTGACACG TTCCT
AACGTCATCTCTCATTCTGTGACACGC TCCT
AACGTCATCTCTCATTCTGTGACACGCA CCT
AACGTCATCTCTCATTCTGTGACACGCAG CT
AACGTCATCTCTCATTCTGTGACACGCAGG T
AACGTCATCTCTCATTCTGTGACACGCAGGG
ACACGCAGGGT ACGTCATCTCTCATTCTGTG

k -mer view


GTCGTTCTTCCTAACGTCATCTCTCATTCTG
TCGTTCTTCCTAACGTCATCTCTCATTCTGT
CGTTCTTCCTAACGTCATCTCTCATTCTGTG
GTTCTTCCTAACGTCATCTCTCATTCTGTGA
TTCTTCCTAACGTCATCTCTCATTCTGTGAC
TCTTCCTAACGTCATCTCTCATTCTGTGACA
CTTCCTAACGTCATCTCTCATTCTGTGACAC
TTCCTAACGTCATCTCTCATTCTGTGACACG
TCCTAACGTCATCTCTCATTCTGTGACACGC
CCTAACGTCATCTCTCATTCTGTGACACGCA
CTAACGTCATCTCTCATTCTGTGACACGCAG
TAACGTCATCTCTCATTCTGTGACACGCAGG
AACGTCATCTCTCATTCTGTGACACGCAGGG
ACGTCATCTCTCATTCTGTGACACGCAGGGT

NECKLACES OF CONSECUTIVE k -MERS SHARE LONG PREFIXES

necklace view

AACGTCATCTCTCATTCTG GTCGTTCTTCCT
AACGTCATCTCTCATTCTGT TCGTTCTTCCT
AACGTCATCTCTCATTCTGTG CGTTCTTCCT
AACGTCATCTCTCATTCTGTGA GTTCTTCCT
AACGTCATCTCTCATTCTGTGAC TTCTTCCT
AACGTCATCTCTCATTCTGTGACA TCTTCCT
AACGTCATCTCTCATTCTGTGACAC CTCCT
AACGTCATCTCTCATTCTGTGACACG TTCCT
AACGTCATCTCTCATTCTGTGACACGC TCCT
AACGTCATCTCTCATTCTGTGACACGCA CCT
AACGTCATCTCTCATTCTGTGACACGCAG CT
AACGTCATCTCTCATTCTGTGACACGCAGG T
AACGTCATCTCTCATTCTGTGACACGCAGGG
ACACGCAGGGT ACGTCATCTCTCATTCTGTG

Size of common prefix
between necklaces of successive k -mers ($k = 31$)



QUICKLY COMPUTING STREAMS OF NECKLACES

Basic approach: compute every cyclic rotation and select the smallest in $\mathcal{O}(k)$.
→ $\mathcal{O}(nk)$ for n necklaces

Better: amortize the computation for consecutive k -mers.

Key observation

If $\langle x \rangle$ does not start at one of the $m - 1$ last positions of x ,
its prefix of size m is the smallest factor of size m in x .

Good news: we can keep track of the smallest factors of size m in $\mathcal{O}(1)$ amortized time using a monotone queue.

m

A	T	A	A	C	G	T	C
T	A	A	C	G	T	C	A
A	A	C	G	T	C	A	T
A	C	G	T	C	A	T	A
C	G	T	C	A	T	A	A
G	T	C	A	T	A	A	C
T	C	A	T	A	A	C	G
C	A	T	A	A	C	G	T

QUICKLY COMPUTING STREAMS OF NECKLACES

Faster necklace computation


Only consider the cyclic rotations that start:

- at one of the smallest factors of size m
- at one of the $m - 1$ last positions

Useful property [Zheng et al. 20]

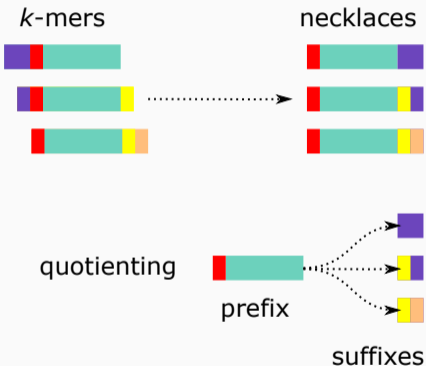
Assuming $m = \Omega(\log k)$, the probability that a k -mer contains **duplicate m -mers** is $o(1/k)$.

By choosing $m = \Theta(\log k)$,
the smallest factor of size m is unique w.h.p.
 $\rightarrow \mathcal{O}(nm) = \mathcal{O}(n \log k)$ for n necklaces (on avg)



DESIGN OF THE DATA STRUCTURE

QUOTIENTING THE PREFIXES OF NECKLACES



Quotienting:

- avoids redundancy
- groups consecutive necklaces


Under the hood:

- store the prefixes in a dynamic bitvector supporting rank/select [Marchini & Vigna 20, Pibiri & Kanda 21]
- associate suffix buckets using a tiered vector (for fast dynamic insertions) [Bille et al. 17]

SCALING THE SUFFIX BUCKETS


The structure of the buckets changes dynamically as we add/remove k -mers

- for small buckets: packed vector, linear search
- for large buckets: trie, logarithmic search




LAYOUT OF CBL'S DATA STRUCTURE

1. compute $\langle x \rangle$
 2. split $\langle x \rangle$ as $q || r$
 3. query r in the bucket of q
- faster for consecutive k -mers



plain necklace vector



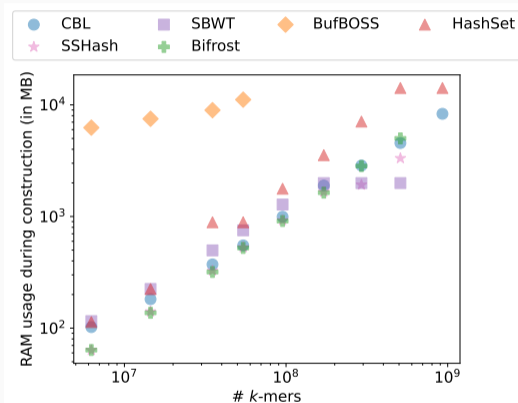
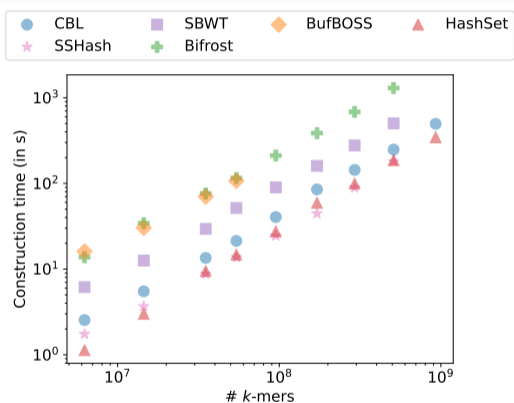
CBL's data-structure

COMPARISON TO SOME EXISTING TOOLS

category	data structure	membership	insert	delete	$\cup \cap \setminus$
BWT	FM-index	✓	✗	✗	✗
—	SBWT	✓	✗	✗	✗
—	dynamic BOSS	✓	✓	✓	✗
hashing	SSHash	✓	✗	✗	✗
—	Bifrost	✓	✓	✗	✗
—	Brisk	✓	✓	✗	✗
—	Bloom filter	approx	✓	✗	union
—	Quotient filter	approx*	✓	✗	union
other	Conway-Bromage	✓	✓	✓	✓
—	CBL	✓	✓	✓	✓

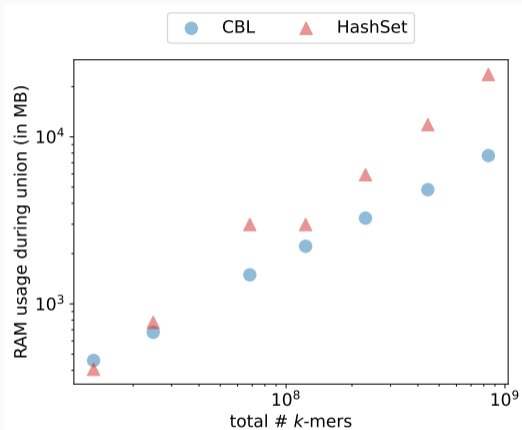
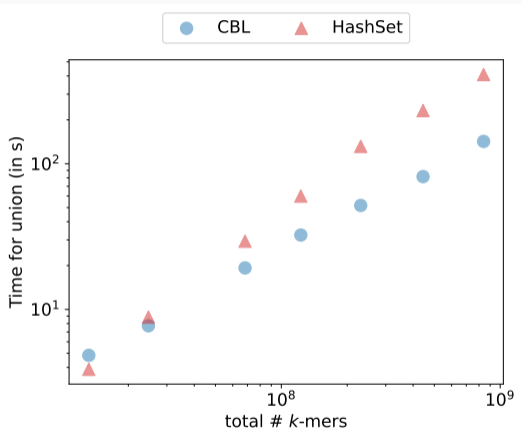
*exact if a PHF is used

TIME/MEMORY USAGE FOR COLLECTIONS OF BACTERIAL GENOMES FROM REFSEQ



TLDR: almost as fast as a hash table, more memory-efficient

MERGING COLLECTIONS OF BACTERIAL GENOMES FROM REFSEQ



TLDR: 4× faster and 3× smaller than a hash table when merging a billion 31-mers

WHAT'S NEXT?

FUTURE STEPS

Improvements of the current structure:

- handle streams of k -mers
- improve buckets' memory usage
(some ideas: smaller single buckets, adaptive radix trie)
- use SIMD for core operations

Extending the structure:

- concurrent version of CBL
(distribute suffix buckets between threads)
- associate data (e.g. count) to each k -mer (\rightarrow map structure)

[Your suggestion here]: let's discuss!

TAKE-HOME MESSAGES

- new dynamic structure based on necklaces
- available as a CLI and a Rust library
- very fast queries, cache efficient
- limited memory usage
- scales for large collections
- versatile operations

Thank you!

Try it here:

github.com/imartayan/CBL










Preprint

(accepted to ISMB)



REFERENCES


-  Alanko, Jarno N, Simon J Puglisi & Jaakko Vuohtoniemi (2022). “**Succinct k-mer sets using subset rank queries on the spectral burrows-wheeler transform**”. In: *bioRxiv*, pp. 2022–05.
-  Bille, Philip et al. (2017). “**Fast dynamic arrays**”. In: *arXiv preprint arXiv:1711.00275*.
-  Conway, Thomas C & Andrew J Bromage (2011). “**Succinct data structures for assembling large genomes**”. In: *Bioinformatics* 27.4, pp. 479–486.
-  Marchini, Stefano & Sebastiano Vigna (2020). “**Compact Fenwick trees for dynamic ranking and selection**”. In: *Software: Practice and Experience* 50.7, pp. 1184–1202.
-  Pibiri, Giulio Ermanno & Shunsuke Kanda (2021). “**Rank/select queries over mutable bitmaps**”. In: *Information Systems* 99, p. 101756.
-  Sawada, Joe & Aaron Williams (2017). “**Practical algorithms to rank necklaces, Lyndon words, and de Bruijn sequences**”. In: *Journal of Discrete Algorithms* 43, pp. 95–110.
-  Zheng, Hongyu, Carl Kingsford & Guillaume Marçais (2020). “**Improved design and analysis of practical minimizers**”. In: *Bioinformatics* 36.Supplement_1, pp. i119–i127.

DENSIFYING THE SPACE OF NECKLACES BY RANKING

The **number of necklaces** of size k on an alphabet with σ letters is

$$N(k) = \frac{1}{k} \sum_{d|k} \varphi\left(\frac{k}{d}\right) \sigma^d \sim \frac{\sigma^k}{k}$$

so only a fraction $\frac{1}{k}$ of the universe is actually used



Ranking: given a necklace $\langle x \rangle$, find i s.t. $\langle x \rangle$ is the i -th smallest necklace of size k

We can compute the rank in $\mathcal{O}(k^2)$ time [Sawada & Williams 17]

Tradeoff: better locality + compression vs $\mathcal{O}(k^2)$ queries