

# *k*-mer data structures in sequence bioinformatics

Rayan Chikhi

Institut Pasteur & CNRS

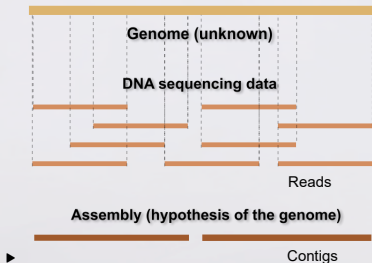
HDR defense, Sep 2021

1. A tale of optimizing the space usage of de Bruijn graphs
2. Minimizer-space de Bruijn graphs

M.	<b>KUCHEROV Gregory</b>	Rapporteur	DR,
Mme	<b>SCHBATH Sophie</b>	Rapportrice	DR,
M.	<b>STOYE Jens</b>	Rapporteur	Professeur,
Mme	<b>CARBONE Alessandra</b>	Examinatrice	Professeur,
M.	<b>GASCUEL Olivier</b>	Examineur	DR,
M.	<b>RICHARD Hugues</b>	Examineur	MCF,
M.	<b>VALLENET David</b>	Examineur	DR,
M.	<b>VINAR Tomas</b>	Examineur	Professeur

# 44 years of genome assembly

- ▶ **1977:** First complete genome assembled (phi X 174)
- ▶ **2003:** Human Genome Project completed
- ▶ **2014:** First \$1,000 genome
- ▶ **2021:** Truly completed (Telomere-2-Telomere)



▶ (Staden 1979) *“With modern fast sequencing techniques and suitable computer programs it is now possible to sequence whole genomes without the need of restriction maps.”*



# Algorithmic pre-history

Contig Editor: -1 SRR030257.787415/2

Cons 2 Qual 0 Insert Edit Modes >> Cutoffs Undo Next Search Commands >> Settings >> Quit Help >>

161080 161090 161100 161110 161120 161130 161140 161150

+17 NC\_012967 GGGTACCAGAACATGGCGGCAAAACAGGAACGCCGGGTTCACGCGCATATCGTTATGGATACGCGTATCGGTATTCTTCAG

+18 \_cer\_sxa\_0\_ GGGTACCAGAACATGGCGGCAAAACAGGAACGCCGGGT

+19 \_cer\_sxa\_60\_ GGGTACCAGAACATGGCGGCAAAACAGGAACGCCGGG

+20 \_cer\_sxa\_123\_ GGGTACCAGAACATGGCGGCAAAACAGGAACGCCGG

+3663 \_cer\_sxa\_190\_ GGGTACCAGAACATGGCGGCAAAACAGGAACGCC

+3664 \_cer\_sxa\_260\_ GGGTACCAGAACATGGCGGCAAAACAGGAACGCC

-33401 SRR030257.8883 ACCAGAACATGGCGGCAAAACAGGAACGCCGGGTGCA

-33402 SRR030257.2989 CCAGAACATGGCGGCAAAACAGGAACGCCGGGTGCA

+33403 SRR030257.1128 CAGAACATGGCGGCAAAACAGGAACGCCGGGTGCACG

+33404 SRR030257.7073 AGAACATGGCGGCAAAACAGGAACGCCGGGTGCACGC

-33405 SRR030257.1204 AGAACATGGCGGCAAAACAGGAACGCCGGGTGCACGC

-33406 SRR030257.2602 GAACATGGCGGCAAAACAGGAACGCCGGGTGCACGCG

-33407 SRR030257.2767 GAACATGGCGGCAAAACAGGAACGCCGGGTGCACGCG

-33408 SRR030257.1413 AACATGGCGGCAAAACAGGAACGCCGGGTGCACGCGC

+33409 SRR030257.1755 ACATGGCGGCAAAACAGGAACGCCGGGTGCACGCGCA

+33410 SRR030257.1463 ACATGGCGGCAAAACAGGAACGCCGGGTGCACGCGCA

+33411 SRR030257.1623 CATGGCGGCAAAACAGGAACGCCGGGTGCACGCGCAT

+33412 SRR030257.6602 TGGCGGCAAAACAGGAACGCCGGGTGG

-33413 SRR030257.1821 GGGCGGCAAAACAGGAACGCCGGGTGCACGC

-33414 SRR030257.2932 GGGCGGCAAAACAGGAACGCCGGG

+33415 SRR030257.5986 GGCAGCAAAACAGGAACGCCGGGTGG

-33416 SRR030257.3729 GGCAGCAAAACAGGAACGCCGGGTGCACGCGCATATC

+33417 SRR030257.3423 GCGGCAAAACAGGAACGCCGGGTGCACGCGCATATCG

+33418 SRR030257.2482 GCGGCAAAACAGGAACGCCGGGTGCACGCGCATATCG

-33419 SRR030257.1401 CGGCAAAACAGGAACGCCGGGTGCACGCGCATATCG

-33420 SRR030257.3565 CGGCAAAACAGGAACGCCGGGTGCACGCGCATATCGT

+33421 SRR030257.1125 CGGCAAAACAGGAACGCCGGGTGCACGCGCATATCGT

-33422 SRR030257.3529 GGCAAAACAGGAACGCCGGGTGCACGCGCATATCGTT

> CONSENSUS -\*- GGGTACCAGAACATGGCGGCAAAACAGGAACGCCGGGTTCACGCGCATATCGTTATGGATACGCGTATCGGTATTCTTCAG /

Tag type:SVEC Direction:- Comment:''

# Algorithmic pre-history



## 1. Assembly using strings

- ▶ Shortest Common Superstring (Kececioğlu, Myers 1993)
- ▶ Greedy algorithms (CAP3 from Huang, Madan 1999)

# Algorithmic pre-history



## 1. Assembly using strings

- ▶ Shortest Common Superstring (Kececioğlu, Myers 1993)
- ▶ Greedy algorithms (CAP3 from Huang, Madan 1999)

## 2. Assembly using graphs: string graphs and de Bruijn Graphs (both from DIMACS'94)

A History of DNA Sequence Assembly, G. Myers, 2016

dBGs widely used across genomics (SPAdes: 13,000 citations; Trinity: 12,000 citations)

# Algorithmic pre-history



## 1. Assembly using strings

- ▶ Shortest Common Superstring (Kececioğlu, Myers 1993)
- ▶ Greedy algorithms (CAP3 from Huang, Madan 1999)

## 2. Assembly using graphs: string graphs and de Bruijn Graphs (both from DIMACS'94)


A History of DNA Sequence Assembly, G. Myers, 2016

dBGs widely used across genomics (SPAdes: 13,000 citations; Trinity: 12,000 citations)

# de Bruijn graph

A **de Bruijn** graph for a fixed integer  $k$ :

1. **Nodes** = all  $k$ -mers (substrings of length  $k$ ) in the reads
2. **Edges** = all exact overlaps of length exactly  $(k - 1)$



Of those reads:

AGC**C**TGA


AGC**A**TGA



# de Bruijn graph

A **de Bruijn** graph for a fixed integer  $k$ :

1. **Nodes** = all  $k$ -mers (substrings of length  $k$ ) in the reads
2. **Edges** = all exact overlaps of length exactly  $(k - 1)$



Of those reads:

AGC**CT**GA

AGC**AT**GA

dBG of *E. coli* reads,  $k=71$ :





Fig: Bandage

# This talk: how we tamed large de Bruijn graphs

E. coli 160,000 genomes  
pangenome mdBG




## The early days (2008-2010)

- Short-read genome assemblers (EULER-SR, Velvet, SOAPdenovo, ABySS)
- **Limited by machine memory** (Most efficient: SOAPdenovo, 120 GB for human)

de Bruijn graph



Hash table




## The early days (2008-2010)

- Short-read genome assemblers (EULER-SR, Velvet, SOAPdenovo, ABySS)
- **Limited by machine memory** (Most efficient: SOAPdenovo, 120 GB for human)

de Bruijn graph



Hash table



```
1 d = dict({
2     'TT': 1,
3     'TG': 1,
4     'AC': 1,
5     'CT': 1
6 })
7 print('CT' in d)
```


# The early days (2008-2010)

- Short-read genome assemblers (EULER-SR, Velvet, SOAPdenovo, ABySS)
- **Limited by machine memory** (Most efficient: SOAPdenovo, 120 GB for human)

## de Bruijn graph



## Hash table



```
1 d = dict({
2     'TT': 1,
3     'TG': 1,
4     'AC': 1,
5     'CT': 1
6 })
7 print('CT' in d)
```

Low contiguity though:

**Table 2.** Summary of bacterial assemblies using 454 reads

Genome	Assembler	No. long contigs	Total length of long contigs (in kb)	N50 (in bases)
<i>S. pneumoniae</i>	EULER-SR	127	2001	32,619
	Newbler	253	2000	11,905
	Repeat graph	136	2091	36,004
<i>E. coli</i>	EULER-SR	199	4277	46,887
	Newbler	141	4531	60,757
	Repeat graph	94	4560	125,693



Table from Chaisson *et al* 2008

## The birth of a line of research (2011)

- Conway & Bromage (2011)
- **Assembly graphs can actually be stored efficiently**

# The birth of a line of research (2011)

- Conway & Bromage (2011)
- **Assembly graphs can actually be stored efficiently**
- Create a large array of  $4^k$  positions (e.g.  $4^{20}$  is a terabit)
- Put **1s** at positions of  $k$ -mers
- Can be **compressed optimally** while supporting queries (Okanohara *et al* 2006)

## de Bruijn graph


A=00	G=10
C=01	T=11
AC=0001	TG=1110
CT=0111	TA=1100



0 1 0 0 0 0 0 1 0 0 0 0 1 0 1 0

0000  
0001  
0010  
0011  
0100  
0101  
0110  
0111  
1000  
1001  
1010  
1011  
1100  
1101  
1110  
1111

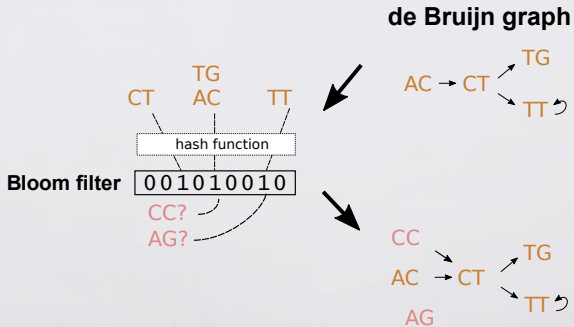
# Is this the end of the research line?





## Beating the lower bound (by inexactness, 2012)

- **2011:** Pell *et al* proposed an inexact dBG representation.
- Bit vector is replaced by a Bloom filter.



# Beating the lower bound (by inexactness, 2012)


- **2011:** Pell *et al* proposed an inexact DBG representation.
- Bit vector is replaced by a Bloom filter.
- **2012:** G. Rizk & I proposed **Minia**
- Same idea, but made the graph exact *where it matters*



Guillaume Rizk



PhD self



# Beating the lower bound (by inexactness, 2012)

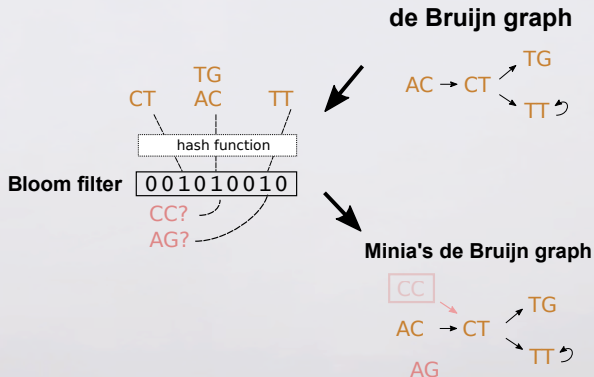
- **2011:** Pell *et al* proposed an inexact DBG representation.
- Bit vector is replaced by a Bloom filter.
- **2012:** G. Rizk & I proposed **Minia**
- Same idea, but made the graph exact *where it matters*



Guillaume Rizk



PhD self



# Beating the lower bound (by inexactness, 2012)

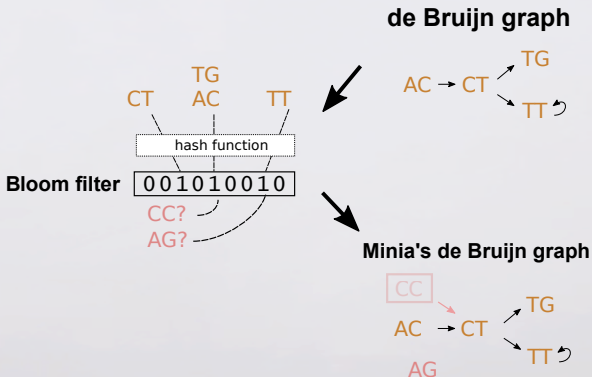


Guillaume Rizk



PhD self

- **2011:** Pell *et al* proposed an inexact DBG representation.
- Bit vector is replaced by a Bloom filter.
- **2012:** G. Rizk & I proposed **Minia**
- Same idea, but made the graph exact *where it matters*
- **Small space!** Beats bit vectors by 2x.



# Beating the lower bound (by inexactness, 2012)

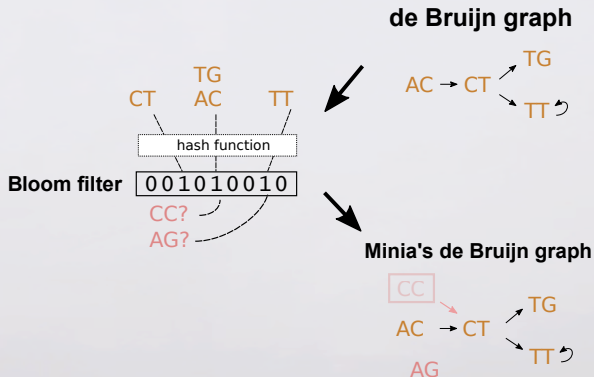


Guillaume Rizk



PhD self

- **2011:** Pell *et al* proposed an inexact DBG representation.
- Bit vector is replaced by a Bloom filter.
- **2012:** G. Rizk & I proposed **Minia**
- Same idea, but made the graph exact *where it matters*
- **Small space!** Beats bit vectors by 2x.
- First assembly of a human genome on a desktop computer.



# Beating the lower bound (by instance specificity, 2012)

- **2012:** Sadakane *et al* proposed the **BOSS** encoding.
- Burrows-Wheeler transform modified to store a set of  $k$ -mers.

	F	Tip	Last		W
A	0	0	1	ACA	C
C	3	1	0	CGA	C
G	8	0	1	CTA	C
T	9	0	0	CAC	G
		0	1	CAC	T
		0	1	GAC	A
		0	0	TAC	A-
		0	1	TAC	T-
		0	1	ACG	\$
		0	1	ACT	A

Fig: MEGAHIT

Fun fact: Minia and BOSS were both introduced at WABI'12

# Beating the lower bound (by instance specificity, 2012)

- **2012:** Sadakane *et al* proposed the **BOSS** encoding.
- Burrows-Wheeler transform modified to store a set of  $k$ -mers.
- **Very small space!** even smaller than Minia.

	F	Tip	Last		W
A	0	0	1	ACA	C
C	3	1	0	\$GA	C
G	8	0	1	CTA	C
T	9	0	0	CAC	G
		0	1	CAC	T
		0	1	GAC	A
		0	0	TAC	A-
		0	1	TAC	T-
		0	1	ACG	\$
		0	1	ACT	A

Fig: MEGAHIT

Fun fact: Minia and BOSS were both introduced at WABI'12

## Beating the lower bound (by instance specificity, 2012)

- **2012:** Sadakane *et al* proposed the **BOSS** encoding.
- Burrows-Wheeler transform modified to store a set of  $k$ -mers.
- **Very small space!** even smaller than Minia.
- But, some limitations (reverse complements, & took years to implement)

	F	Tip	Last		W
A	0	0	1	ACA	C
C	3	1	0	\$GA	C
G	8	0	1	CTA	C
T	9	0	0	CAC	G
		0	1	CAC	T
		0	1	GAC	A
		0	0	TAC	A-
		0	1	TAC	T-
		0	1	ACG	\$
		0	1	ACT	A

Fig: MEGAHIT

Fun fact: Minia and BOSS were both introduced at WABI'12



## New perspective on the topic (2014)

- So, how comes Minia & BOSS beat the lower bound?

---

<sup>1</sup>R Chikhi, A Limasset, S Jackman, JT Simpson, P Medvedev, *On the representation of de Bruijn graphs*, RECOMB'14

## New perspective on the topic (2014)

- So, how comes Minia & BOSS beat the lower bound?
- The lower bound assumed the graph was **exact**.
- Minia only supports **some** operations exactly.

---

<sup>1</sup>R Chikhi, A Limasset, S Jackman, JT Simpson, P Medvedev, *On the representation of de Bruijn graphs*, RECOMB'14

## New perspective on the topic (2014)



Paul Medvedev



Antoine Limasset



Postdoc self


- So, how comes Minia & BOSS beat the lower bound?
- The lower bound assumed the graph was **exact**.
- Minia only supports **some** operations exactly.
- We came up with new lower bounds, i.e.  $\approx 3$  bits instead of 22 bits.<sup>1</sup>
- Open problem: a matching upper bound in the general case
- *Intriguingly fun fact*<sup>2</sup>: BOSS is fully exact (same as bit vector) and yet still beats the lower bound

---

<sup>1</sup>R Chikhi, A Limasset, S Jackman, JT Simpson, P Medvedev, *On the representation of de Bruijn graphs*, RECOMB'14


<sup>2</sup>For the handful of people on Earth who find this fun

# Where are we now? (& my contribs)



# Current works on DBGs (2016-2020)


<i>k</i> -mer counting	Compaction	MPHF based	FM-Index based	General purpose
KMC3	BCALM2*	Pufferfish	DBGFM*	BBHash*
DSK2*	Cuttlefish	BLight	BOSS	Bifrost
Jellyfish2	TwoPaCo	FDBG*	dynamicBOSS	
SPAdes-kmercounter			bufBOSS	




Nowadays:

- Focus is less on space, more on features
- Fast query times, associativity, dynamicity

**Usual hashing**



**Minimal perfect hashing**



RESEARCH-ARTICLE [OPEN ACCESS](#)

## Data Structures to Represent a Set of $k$ -long DNA Sequences

---

**Authors:**  [Rayan Chikhi](#),  [Jan Holub](#),  [Paul Medvedev](#) [Authors Info & Affiliations](#)

---

ACM Computing Surveys, Volume 54, Issue 1 • April 2021 • Article No.: 17, pp 1–22 • <https://doi.org/10.1145/3445967>

---

**Published:** 08 March 2021

## Part 2

# Long reads genome assembly

- Oxford Nanopore, PacBio CLR
  - ▶ 10-1,000 kbp reads, **5-12%** error rate
- PacBio HiFi
  - ▶ 10-25 kbp reads,  $\leq$  **1%** error rate




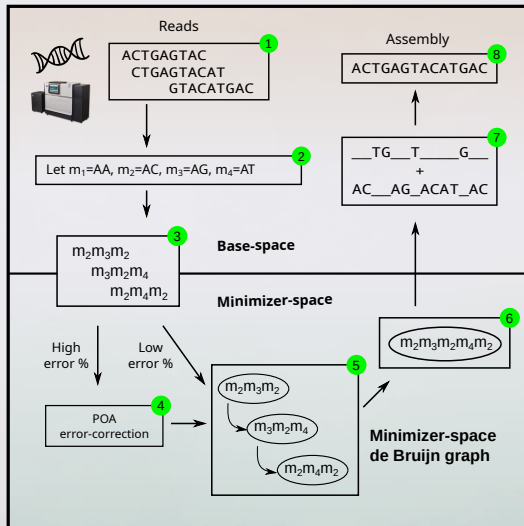
Classical *de Bruijn* graphs not applicable (no long error-free *k*-mers). Instead:

- Overlap graphs (Canu, miniasm, Shasta, Peregrine, ...)
- Fuzzy dBGs (wtDBG2)
- Sparse dBGs: A-Bruijn or minimizers (Flye, MBG)


**Challenge:** Approaches don't scale (high resource usage, slow assembly time)!



# dBGs on long reads: Minimizer-space de Bruijn graphs



Barış Ekim



Bonnie Berger

- Long read human genome assembly on a desktop computer

# Preliminaries: Minimizers

Two kinds:

- **window**. Local: “smallest”  $l$ -mer in a window

AATGACATGATCATGA

AA

AC

AC

- **universe**. Global: set of  $l$ -mers with low hash values

Fixed set of  
universe minimizers




AATGACATGATCATGA

GA

TC

From now on: **universe**.

# This work: stems from three ideas



# Our approach: Minimizers as *tokens* of the alphabet

Classical alphabet:  $\Sigma_{DNA} = \{A, C, T, G\}$

A  $k$ -mer with  $k = 3$ :  $AGT$

**Minimizer alphabet:**

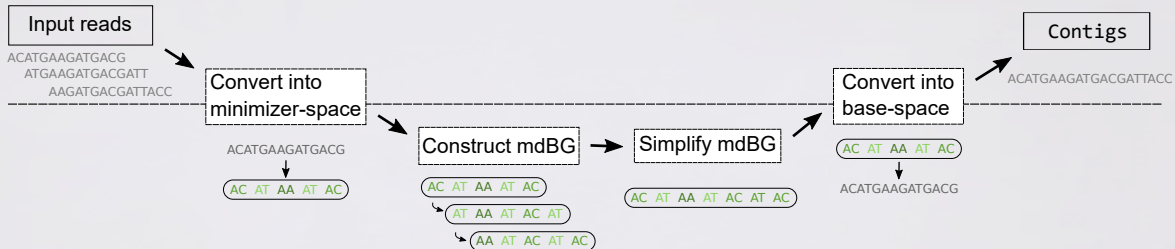
$$\Sigma^\ell = \{\text{all minimizers of length } \ell\} = \{m_1, m_2, m_3, \dots\}$$

where e.g.  $\ell = 2$ ,  $m_1 = AA$ ,  $m_2 = AC$ ,  $m_3 = AG$ ,  $m_4 = AT$

A  $k$ -mer over  $\Sigma^\ell$  (a  $k$ -**min-mer**):  $m_1 m_3 m_2$

# Results: Whole-genome *de novo* assembly

From accurate HiFi (< 1% error-rate) reads



Whole human PacBio HiFi (HG002) 50x coverage:

Tool name	Peregrine	hifiasm	rust-mdbg
Wall-clock time	14h8m	58h41m	<b>10m23s</b>
Memory usage	188 GB	195 GB	<b>10 GB</b>
# contigs	8109	431	805
NG50 (Mbp)	18.2	88.0	16.1
Genome fraction	97.0%	94.2%	95.5%

# Results: Metagenome assembly

## Zymo D6331 mock metagenome HiFi

Species	Abundance	hifi asm-m	rust-mdbg
<i>A. muciniphila</i>	1.36%	100.000%	100.000%
<i>B. fragilis</i>	13.13%	99.994%	99.997%
<i>B. adolescentis</i>	1.34%	100.000%	99.730%
<i>C. albican</i>	1.61%	67.832%	39.821%
<i>C. difficile</i>	1.83%	99.996%	99.978%
<i>C. perfringens</i>	0.00%	0.005%	0.005%
<i>E. faecalis</i>	0.00%	0.006%	0.006%
<i>E. coli B1109</i>	8.44%	100.000%	97.918%
<i>E. coli b2207</i>	8.32%	100.000%	98.663%
<i>E. coli B3008</i>	8.25%	100.000%	99.558%
<i>E. coli B766</i>	7.83%	96.913%	96.270%

Species	Abundance	hifi asm-m	rust-mdbg
<i>E. coli JM109</i>	8.37%	100.000%	97.852%
<i>F. prausnitzii</i>	14.39%	100.000%	100.000%
<i>F. nucleatum</i>	3.78%	100.000%	99.960%
<i>L. fermentum</i>	0.86%	100.000%	100.000%
<i>M. smithii</i>	0.04%	99.840%	87.175%
<i>P. corporis</i>	5.37%	99.561%	99.561%
<i>R. hominis</i>	3.88%	100.000%	100.000%
<i>S. cerevisiae</i>	0.18%	69.522%	39.556%
<i>S. enterica</i>	0.02%	6.232%	4.619%
<i>V. rogosae</i>	11.02%	100.00%	100.000%


	hifi asm-m	rust-mdbg
<b>Running time</b>	34h29m	<b>55s</b>
<b>Memory usage</b>	83 GB	<b>0.9 GB</b>

# For > 1% error rates: Minimizer-space POA error correction

(base-space POA: Lee *et al*, 2002)

ACGGATTCACGGAA    ACAGATTCCGGTA  
 $\overline{m_1}$     $\overline{m_2}$     $\overline{m_1}$     $\overline{m_3}$     $\overline{m_1}$   $\overline{m_4}$   $\overline{m_2}$

ACGGATTCCGGAAT  
 $\overline{m_1}$     $\overline{m_2}$     $\overline{m_3}$   $\overline{m_2}$

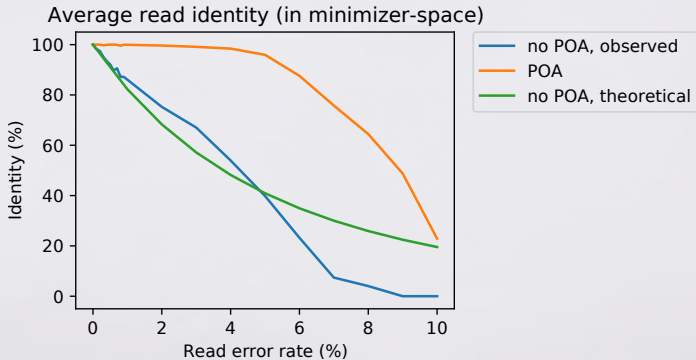


Final consensus

$\overline{m_1 m_2 m_3}$

# For > 1% error rates: Minimizer-space POA error correction

(base-space POA: Lee *et al*, 2002)




So, not quite ready for Nanopore data ( $\geq 5\%$ ).



# Results: Pangenome graph of 661,405 bacterial genomes

Largest 5  
connected  
components:



# Biological results: Querying AMR genes

AMR genes  
database



1,279 genes  
(AMRFinderPlus)


Graph query

ACATGAAGATGACGATTACC

Convert to minimizer-space


ACATAAATAC AT AC

Query each k-min-mer



Min-space query coverage : 2/3

Retrieval of AMR genes



## Part 3 (short)

# K and U problems

Known problems



Img: freepic

Can outline research plan.

Unknown problems



Img: freepic

Previously thought impossible.

# K and U problems

## Known problems



Img: freepic

Can outline research plan.

E.g.:

- mDBG
- BCALM2
- REINDEER

## Unknown problems



Img: freepic

Previously thought impossible.

E.g.:

- ▶ ▶ Minia
- ▶ BOSS
- ▶ pugz

Conclusion

# Future directions

In the dBG area:

- Representations of **multiple samples**: REINDEER, BFT, HowDeSBT, MetaGraph, etc.. (Marchet *et al* review in Genome Res'20)
- Efficient storage of **abundances**: Italiano *et al*; Shibuya & Kucherov, ...
- Best adaptation to **long reads**: wtdbg2, mdBG, Flye, ...
- **Disk** compression: SPSS, Simplitigs, ...
- A standard file format: [github.com/Kmer-File-Format](https://github.com/Kmer-File-Format)

And advising team projects:

- Metagenomics strain assembly
- Ancient DNA decontamination
- Structural variants detection
- Sequence transformations

## SeqBio Group @ Institut Pasteur



Y. Dufresne, R. Vicedomini, L. Denti, T. Lemane, C. Duitama, L. Blassel

And former students: Camille Marchet, Pierre Marijon, Maël Kerbiriou

And all my current and previous collaborators: I had a wonderful list but it was too long to fit inside this slide <3

Funding: RiSE Pangaia, ITN Alpaca, ANR Inception, ANR Prairie, ANR Transipedia, ANR SeqDigger





**Lex Nederbragt**

@lexnederbragt

En réponse à [@ctitusbrown](#)

“Finding your way in life is like finding the genome in a De Bruijn graph: it is very easy to find *\*a\** path, very hard to find *\*the\** path”.

Thank you all for your attention!

M.	<b>KUCHEROV Gregory</b>	Rapporteur	DR,
Mme	<b>SCHBATH Sophie</b>	Rapportrice	DR,
M.	<b>STOYE Jens</b>	Rapporteur	Professeur,
Mme	<b>CARBONE Alessandra</b>	Examinatrice	Professeur,
M.	<b>GASCUEL Olivier</b>	Examineur	DR,
M.	<b>RICHARD Hugues</b>	Examineur	MCF,
M.	<b>VALLENET David</b>	Examineur	DR,
M.	<b>VINAR Tomas</b>	Examineur	Professeur

## Bit vector optimality

- A de Bruijn graph **only needs to records the nodes**.
- **Bijection** between **sets of nodes** and **binary vectors** of length  $4^k$ .
- How many different bit vectors of size  $4^k$  and  $n$  1's?

$$\binom{4^k}{n}$$

- Thus, **minimal number of bits** to store a dBG:

$$\log_2\left(\binom{4^k}{n}\right)$$

## Bit vector optimality

- A de Bruijn graph **only needs to records the nodes**.
- **Bijection** between **sets of nodes** and **binary vectors** of length  $4^k$ .
- How many different bit vectors of size  $4^k$  and  $n$  1's?

$$\binom{4^k}{n}$$

- Thus, **minimal number of bits** to store a dBG:

$$\log_2\left(\binom{4^k}{n}\right)$$

- A compressed bit vector achieves this optimal space.

## Bit vector optimality

- A de Bruijn graph **only needs to records the nodes**.
- **Bijection** between **sets of nodes** and **binary vectors** of length  $4^k$ .
- How many different bit vectors of size  $4^k$  and  $n$  1's?

$$\binom{4^k}{n}$$

- Thus, **minimal number of bits** to store a dBG:

$$\log_2\left(\binom{4^k}{n}\right)$$


- A compressed bit vector achieves this optimal space.
- (This is much smaller than  $O(kn)$ , the hash table storage)

## Caveats

- Only a subset of approaches were presented
- Ignored query times
- Ignored associated info (e.g.  $k$ -mer abundances)
- Ignored analysis environment (error-correction, assembly algorithms)
- Ignored multi- $k$
- Ignored reverse-complements
- Ignored the rest of the bioinformatics field, biology, etc..

## Recommended readings

### Modeling biological problems in computer science: a case study in genome assembly


Paul Medvedev 

&

The theoretical analysis of sequencing bioinformatics algorithms  
(DRAFT)

Paul Medvedev<sup>1,2,3</sup>

### What do Eulerian and Hamiltonian cycles have to do with genome assembly?

Paul Medvedev , Mihai Pop

Published: May 20, 2021 • <https://doi.org/10.1371/journal.pcbi.1008928>