

Molecular Evolution and **Bioinformatics**



Reference based sequence mapping

Literature

How to map billions of short reads onto genomes?

Trapnell and Salzberg *Nature Biotechnology* 27, 455 - 457 (2009)

doi:10.1038/nbt0509-455


Ultrafast and memory-efficient alignment of short DNA sequences to the human genome.

Langmead et al. *Genome Biology* **10**:R25 (2009)


doi:10.1186/gb-2009-10-3-r25

What do we need short read mapping for?

Assembly of whole genome shotgun sequencing data

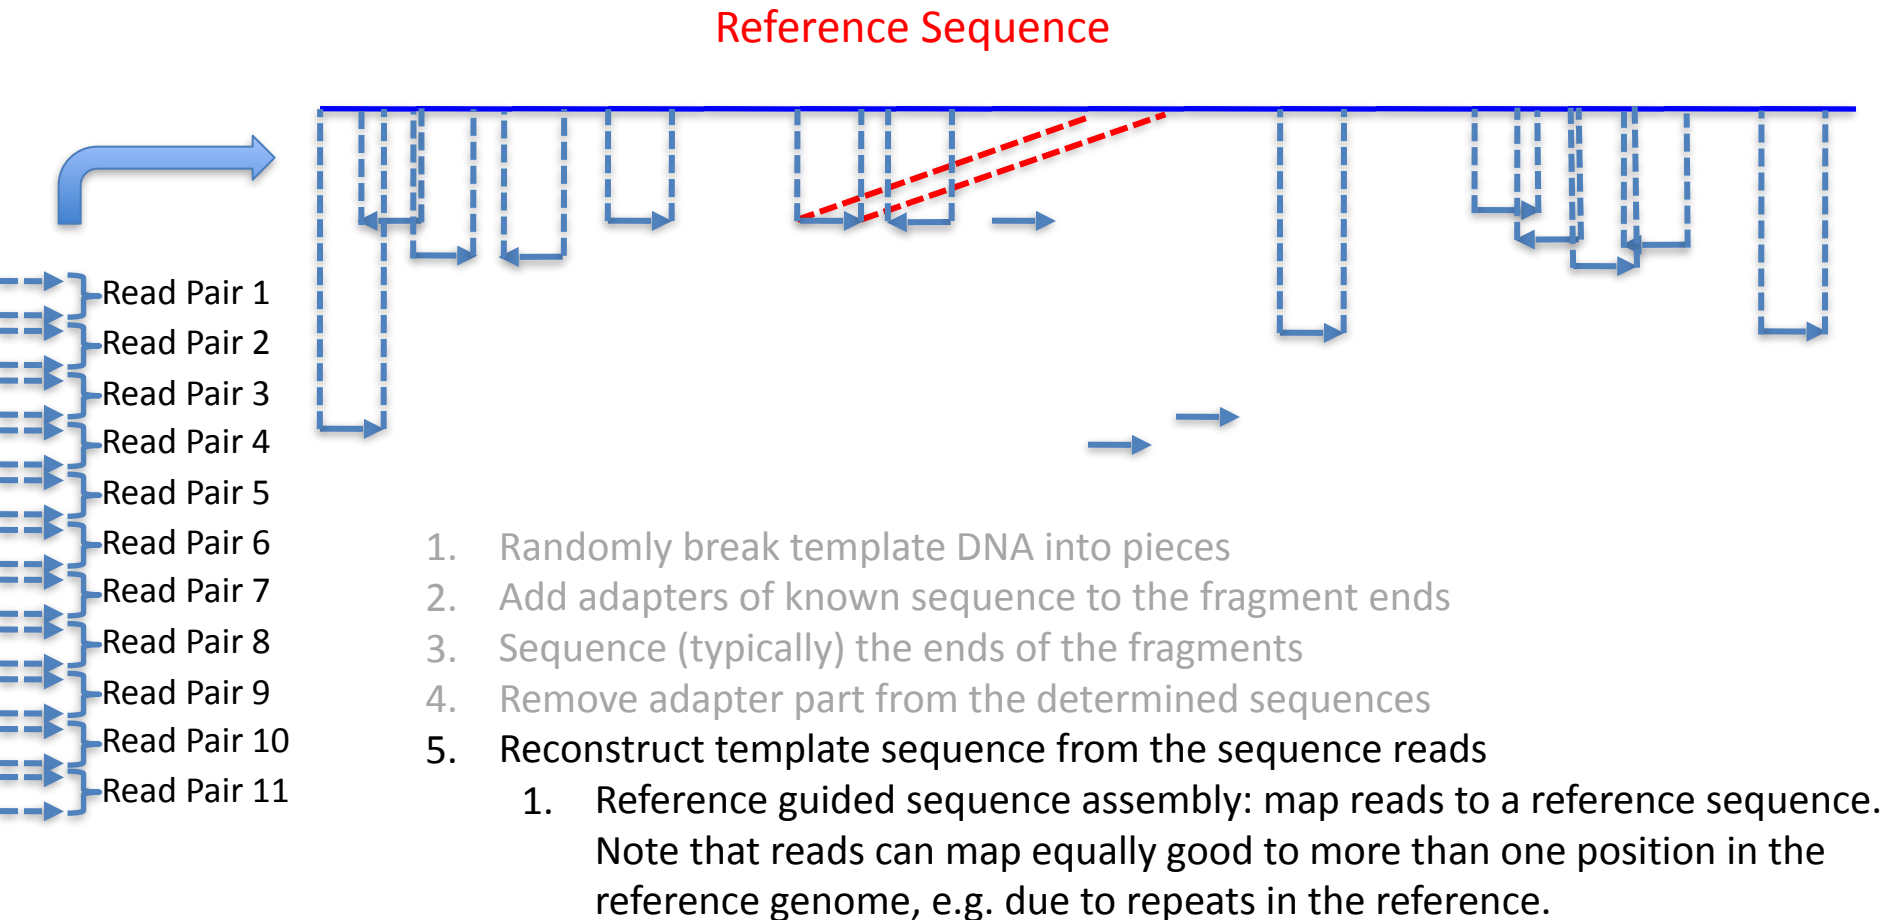


Reconstruct template



1. Randomly break template DNA into pieces
2. Add adapters of known sequence to the fragment ends
3. Sequence (typically) the ends of the fragments
4. Identify and remove adapter part from the determined sequences
5. Reconstruct template sequence from the sequence reads

Shotgun sequencing and **reference guided** sequence assembly

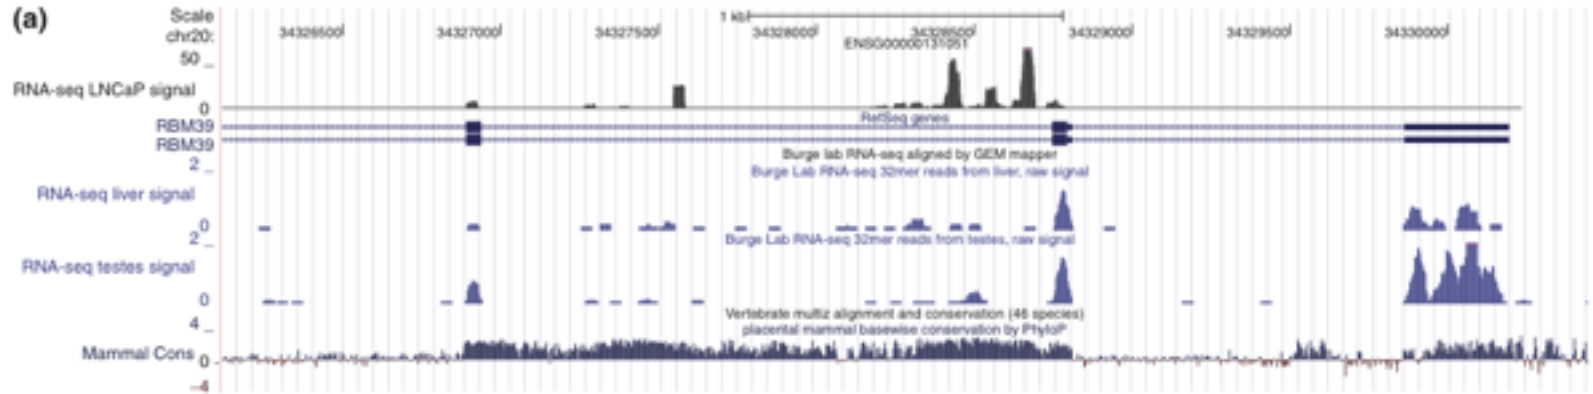


— Reference sequence

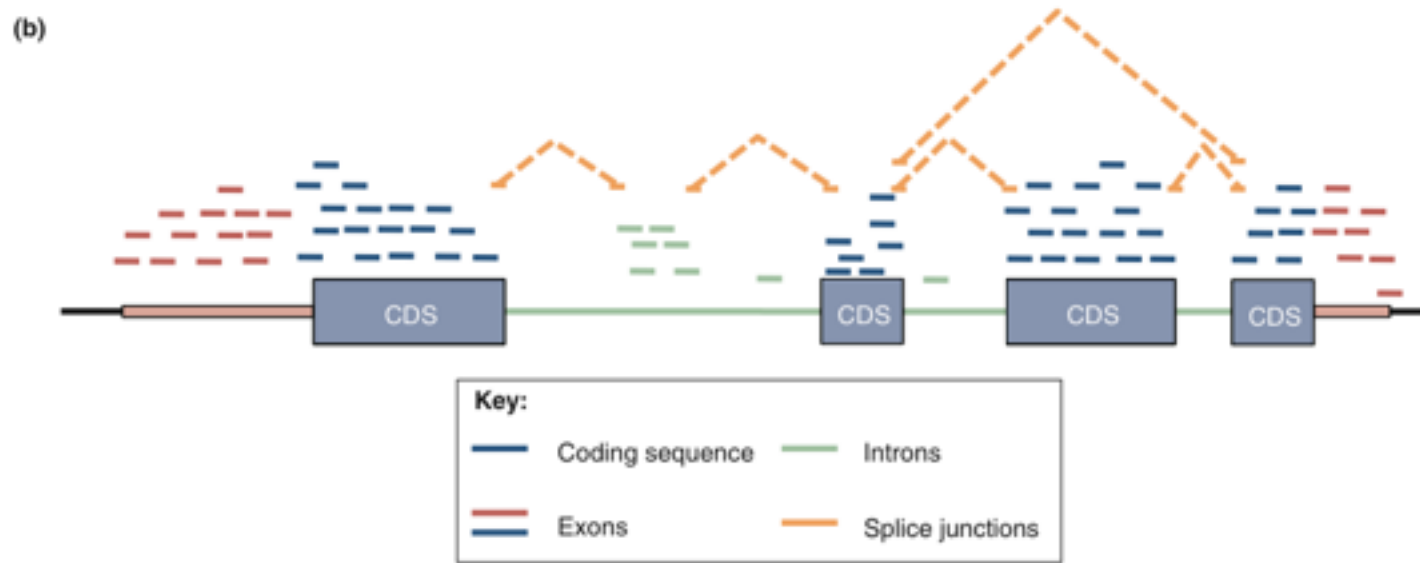

a) e.g. genome of a different individual from the same species to study species diversity

b) e.g. genome of a closely related species

Further short read applications: Mapping of RNA seq



RNA-seq mapping helps building hypotheses genes concerning gene structure



Further short read applications

- Genotyping

Goal: identify variations

```

...CCATAG      TATGCGCCC      CGGA AATTT  CGGTATAC  GGTATAC...
...CCAT      CTATATGCG      TCGG A AATT  CGGTATAC
...CCAT GGCTATATG      CTATCGG AAA  GCGGTATA
...CCA AGGCTATAT      CCTATCGG A  TTGCGGTA  C...
...CCA AGGCTATAT      GCCCTATCG      TTTGCGGT  C...
...CC  AGGCTATAT      GCCCTATCG  AAATTTGC  ATAC...
...CC  TAGGCTATA GCGCCCTA  AAATTTGC  GTATAC...
...CCATAGGCTATATGCGCCCTATCGGCAATTTGCGGTATAC...
    
```

- ChIP-seq, CLIP-Seq, Methyl-seq

Goal: classify, measure significant peaks

```

          GAAATTTGC
          GGAAATTTG
          CGGAAATTT
          CGGAAATTT
          TCGGAAATT
          CTATCGGAAA
          CCTATCGGA      TTTGCGGT
GCCCTATCG      AAATTTGC
GCCCTATCG      AAATTTGC      ATAC...
...CCATAGGCTATATGCGCCCTATCGGCAATTTGCGGTATAC...
    
```

Challenges

- mapping millions/billions of reads to a large genome is hard
 - how quickly can we map the reads to the genome?
 - how do we deal with multiple mapping positions?
 - how do we deal with sequencing errors and genetic divergence/diversity
 - how do we deal with reads that span intron-exon boundaries?

Short Read Alignment

- Given a reference and a set of reads, report at least one “good” local alignment for each read if one exists
 - Approximate answer to: from where in genome did the read originate?

What is “good”?

- Fewer mismatches is better
- Failing to align a low-quality base is better than failing to align a high-quality base

...TGATCATA... better than ...TGATCATA...
| | | | | | | |
GATCAA GAGAAT

...TGATATTA... better than ...TGATcaTA...
| | | | | | | |
GATcaT GTACAT

Finding mapping positions of reads in genomes is, in principle, a pattern matching problem. We search for the occurrence of short words in long sentences.

Finding mapping positions is, in principle, very easy!

Searchstring: AATGAGACATGAA

Query1: CATG

Query2: ATGT

Finding mapping positions is, in principle, very easy. Just slide the word along the sequence and stop when either end of sequence is reached or mapping position is found


Searchstring: AATGAGACATGAA

Query1: CATG

Query2: ATGT



Naïve approach

Searchstring: 1234567891234
AATGAGACATGAA
CATG 

Finding mapping positions is, in principle, very easy. Just slide the word along the sequence and stop when either end of sequence is reached or mapping position is found

Searchstring: AATGAGACATGAA

Query1: CATG

Query2: ATGT



Naïve approach

Searchstring: 1234567891234
AATGAGACATGAA
CATG
✗+++

Finding mapping positions is, in principle, very easy. Just slide the word along the sequence and stop when either end of sequence is reached or mapping position is found

Searchstring: AATGAGACATGAA

Query1: CATG

Query2: ATGT



Naïve approach

Searchstring: 1234567891234
AATGAGACATGAA
CATG
XXXXXX

Finding mapping positions is, in principle, very easy. Just slide the word along the sequence and stop when either end of sequence is reached or mapping position is found

Searchstring: AATGAGACATGAA

Query1: CATG

Query2: ATGT



Naïve approach

Searchstring: 1234567891234
AATGAGACATGAA
CATG
xxx+x

Finding mapping positions is, in principle, very easy. Just slide the word along the sequence and stop when either end of sequence is reached or mapping position is found

Searchstring: AATGAGACATGAA

Query1: CATG

Query2: ATGT



Naïve approach

Searchstring: AATGAGACATGAA

1234567891234

CATG

++++

Query1 maps to position 8-12 in searchstring



Full match found,
Output result

Finding mapping positions is, in principle, very easy. Just slide the word along the sequence and stop when either end of sequence is reached or mapping position is found

Searchstring: AATGAGACATGAA

Query1: CATG

Query2: ATGT



Naïve approach

Searchstring: 1234567891234
AATGAGACATGAA
TTGT

XXXXX

At most $n-k$ comparisons, with n is the length of the search string, and k is the query length (read length).

This is not feasible for short read mapping.

Indexing speeds up searches

INDEX					
Adams, Joe	80-81	Collins, Clarence	26	Foster, John Wesley	78
Alvarado, Hiram O.	81	Conover, Benjamin Edward	45	Foster, T. W.	79
Arnold, Dan and Bettina	43-42	Conover, E. J.	85, 138	Frio County Centennial	3
Arnold, George and		Conover, Freddie Marvis	85, 138	Frio County Historical	
Agatha	82	Conover, Fred N.	45-47	Commissions	4
Arnold, Henry and Clara	18-19	Conover, George W.	47	Frio County History	4
Assembly of God Church		Conover, George Washington	47	Frio County Markers	4-5
(Dilley)	17	Conover, Mae D.	47	Frio County Tax Capital	5-7
Assembly of God Church		Conover, Minnie	47	Frio Pioneer Jail Museum	6
(Prescott)	38	Conover, William O.	47	Frio Public Library	7-8
Asst. Forner J.	19	County, Roosevelt and Loh	68-70	Fudge, Albert Lester	73-80
Asst. James Ross	19-20	Cowden, George	70-71	Gerr, E. F.	18-11
Asst. Robert F. and		Cowley, W. B.	71-72	Goodman, Nancy Johnson	83
Florida	20	Cox, Joseph	72	Goss, L. E.	83
Beall, Dr. J. E.	42-43	Crawford, V. T. and Mary	72, 129	GR, Hugh E.	50-51
Benson, John	83	Crocker, Missie	72	Gosner, Whitford	10
Berry, James E.	63-64	Crossen, James	147	Haacke, Henry J.	40-41
Berry, Mrs. James E.	64	Cude, Frank	39	Harrell, William D.	11
Berry, News Ward	64	Cude, Lowell	39	Harris, W. A. and Nancy	158
Bigler Baptist Church	9	Cude, Willis Franklin	73-73	Harris, J. Will	31
Bigler Church of		Dalton, Harry and Evelyn	73-74	Harris and Hinder	81
Christ	9		140	Hayes, Wynn Washam	41-42
Bigler Methodist		Danchak, John Michael	74	Hessau, William F.	50
Church	9	Danchak, Joseph Franklin	74-75	Herron, James	82
Bikarta, August and		Davenport, Mary	75-76, 126	Higgins, George J.	42-43
Tracie	64	DeVillio, Luther	26-27	Higgins, Grady and Ruth	83
Bikarta, Joseph and		DeWoods, Almon	26-27	Higgins, James	43
Hortense E.	64-65	DeWoods, T. V.	27	Hiles, Daniel L.	10-32
Bickler, James		Dillard, Nathan	27	Hill, W. S.	43-45
Harrison	65	Dilley Church of Christ	11	Hiler, James H.	45
Biss, William	45-46	Dilley Presbyterian Church	18	Hiler, Lee Reavis	45-46
Biswell, Miss Bird	46	Douglas, Glenn A. and		Hoopas, Margot	86
Boyd, William A.	5-10	Minnie L.	36	Holland, Albert Goss	86-87
Breen, John George	45-46	Dubow, Ruby K.	46	Holmes, Joseph M.	11-12
Brenstein, Maria H.	46-47	Dunn, Claude Bernard	27-28	Howard, Dr. and Mrs. E. M.	87, 149
Brown, Bernard	20-21, 117	Dunn, Hubert	28	Howard, Earl Winfield	47-48
Brown, Charles A.	21	Dunn, John Anthony	28	Howerton, James	12-13
Brown, Frank S.	21	Dunn, Oscar James	28	Inmaculate Heart of Mary	
Brown, James Gibson	21-22	Dunn, Patrick	28	Parish	59
Brown, Milton B.	22	Edwards, Elias H. Crain	48	Johlin, William	88
Brown, Paul P.	22	Edwards, Lex J. Wangrober	48	Johnson, Quill	48-49
Buby, Ales and Gladie	22	Elkins, J. W.	28-29	Joos, Charles Calvin	48
Buby, Claude L.	22-23	Ellis, Houston H.	29	Joos, Floyd C.	49
Buby, Robert	147	Ellis, S. H.	29-30	Joos, O. B. and Michaela	49-50
Buby, Roy and Lucile	23	Ellis, W. W.	30, 141	Joos, James Marion	50
Campbell, Alta	47-48	Ferguson, Claid O.	45-49	Joos, Jeremiah Deeman	50
Casper, Lois	48	Ferguson, John E., Jr. and		Joos, Michael Joe (Mike)	96, 100
Cassick, Rev. Walter	23-24	Mildred	49	Justin, Sam H.	90-91
Carroll, Berny R.	24	Ferguson, John E., Sr.	49-50	Karnes, Blanche Crawford	31
Carroll, Charles L.	24	Fields, John E.	34	Kemper, John	51-52
Carroll, James Henry	24	First Baptist Church		Keshish, Rev. James Floyd	42-43
(Dilley)	17	(Dilley)	17	King, J. J. and Mabel	43-44, 143
Geneva	24	First Baptist Church		Klopp, Albert and Lucille	34, 143
(Prescott)	38	(Prescott)	38	Lee, Charles Edward	53
Carroll, Louis Joseph, Jr.	24-25	First Christmas Church		Leis, Owen E.	50-51
Carroll, Louis Joseph, Sr.	24-25	(Prescott)	38	Lindholm, John Nelson	35-34
Carroll, Paul Eugene	25	First United Methodist		Lindholm, Mady	94
Carroll, Wesley Raymond	25	Church (Dilley)	17-18	Lippard, David	94-95
Carroll, Wesley Robert	25	First United Methodist		Liska, Reix	30
Church (Prescott)	38	Church (Prescott)	38-40	Littman, Sam and Osa	51
Cassidy Family	25-26	Fisk, R. D.	74-77	Long, R. O.	95-96
Church of Christ (Prescott)	38-39	Flores, Ted	77	Low, Judge Marcellus	96
Chauswitz, Martin	48	Foreman, Berne and Beata	77-78	Lynn, Arthur E.	96
Chisman, Robert	68-69	Foster, Lorrna	78-79, 143	Malone, John J.	97, 148
Collins, John W. (Bill)	49				

Indexing speeds up searches

Searchstring: CATGAGACATGAA

Query1: GAGA

Query2: CATG

Query3: ATGT

- 1) Decide on a word length k , e.g., $k=3$
- 2) Build hash table from search string, storing every word occurring in S together with its start position.
- 3) Process query and search for each word occurring in $Q1$ whether it is in the hash table.
- 4) Repeat for $Q2$.

Kmer	Position
CAT	1,8
ATG	2,9
TGA	3,10
GAG	4
AGA	5
GAC	6
ACA	7
GAA	11

Kmer	Position
GAG	1
AGA	2

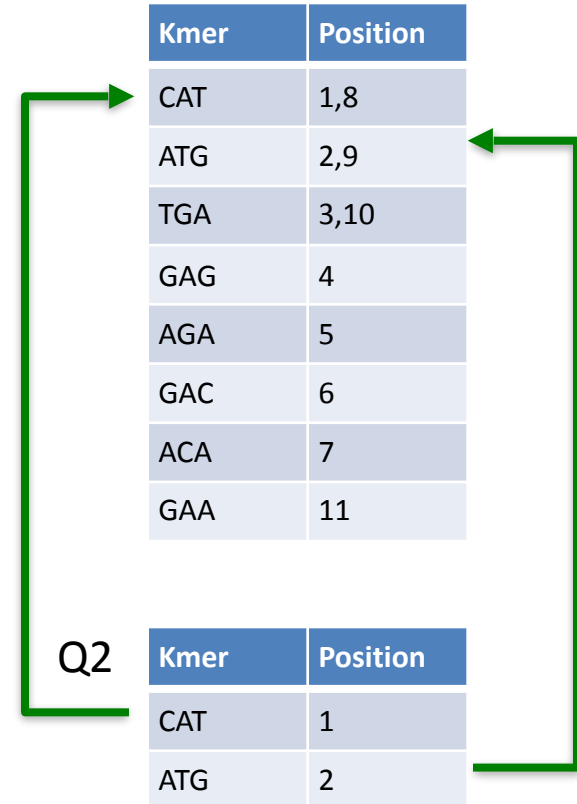
Two lookups are sufficient to find $Q1$ in S

Main differences between mapping approaches

1) How does the mapper decide when the query maps multiple times to the sequence?

Mapping position of Q2: 1-4 **AND** 8-11

Relevant e.g. for RNA seq quantifying gene expression and for Chip-Seq as you must not consider a read more than once.

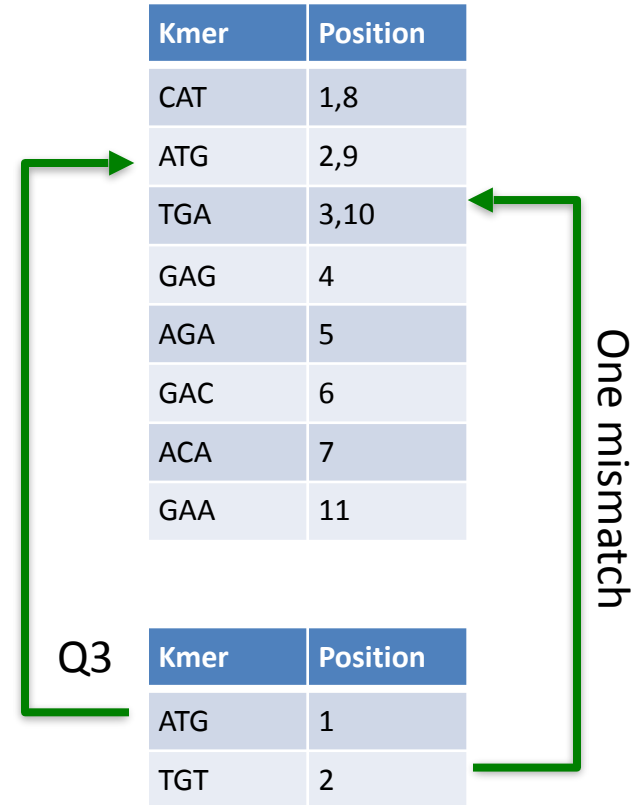


Main differences between mapping approaches

2) How does the mapper deal with queries that 'almost' match the reference?

Q3 matches the reference with one mismatch

Relevant for sensitivity and specificity of the mapping. Allowing more mismatches increases sensitivity (consider sequencing error and genetic diversity) but decreases specificity (more false positives).




The 2nd lookup indicates that Q3 is 'almost' in S


Main differences between mapping approaches

3) What kind of index does the mapper use?


Relevant for speed and memory footprint of the mapper



Suffix tree



Suffix array



Seed hash tables

Many variants, incl. spaced seeds

\$BANANA
 A\$BANAN
 ANAS\$BAN
 ANANAS\$B
 BANANAS\$
 NAS\$BANA
 NANAS\$BA

Burrows Wheeler Transformation


The “traditional way”: Hash tables

- Used by MAQ, Eland, SOAP, SHRiMP, ZOOM, partially by Mosaik



The use of 'hashing' in exact pattern search (Rabin-Karp; $O(n+m)$)

Approach:

- 1) Use a 'hash-function' to transform pattern P into a numerical hash value h_p .



- 1) Search the text T starting from left for words of length $|P|$ having the same hash value as P .




- 1) Given a word K with $h_k = h_p$ was found, perform an exact string comparison to verify that $K == P$. (Note, the projection of words with length $|P|$ in the space of hash values is **not** injective (linkseindeutig!).)

*JUST AN EXAMPLE: $v(A)=1, v(C)=2, v(G)=3, v(T)=4; mod=3$

‘hashing’ in combination with hash tables help to reduce the average time complexity of the pattern search to $O(1)$, i.e. constant in time*

Idea: Speed up pattern search by creating look-up tables storing the hash values

- 1) Search the text T starting from left for words of length k and compute their hash value h_k



- 1) Store the hash values together with the starting point of the corresponding words in a hash table. Note, a hash table is nothing but a special way of indexing your data, just like a phone book. This will provide direct access to your potential matches in the pattern search once you know the hash value of P .

Hash value	Position in string
0	10,11,20,25,26
1	1,2,6,7,8,12,15,18,19,21,22,23,24,27,28
2	3,4,5,9,13,14,16,17,29

MAQ* uses spaced seeds for mapping

Example: read1 - gatgtgacatacctgttctactgaggct

Build **six** hash tables (templates) for the reads (only first 28 bp are considered) **only** from the colored nucleotides

6 'Templates'

gatgtgacatacctgttctactgaggct	HASH → 2057673064
gatgtgacatacctgttctactgaggct	→ 3178370917
gatgtgacatacctgttctactgaggct	→ 773088662
gatgtgacatacctgttctactgaggct	→ 1856750201
gatgtgacatacctgttctactgaggct	→ 2510061809
gatgtgacatacctgttctactgaggct	→ 119777054

hash value for template 1: 832589471

GENOME
tcaagttatccagaaaatgctatTTTTCCGaaaatgaaatctaaaatagtaactcaagtgaacattgtcaggtgtgtaaggaaggaaaatatgaagggctaccacaaaaccacaaaacacggaaaactccaattcctaagtttccgggacacactattcatatagatataatTTCTACAGACAAAAACGGTACTTACGGCAATTCACAAATTTTCAAAACTTGCGAAAGCAAAAAATAAAAAAATTCAAAATCAATAGAAGACATAAGAAAACCTTACATGACATCTTATTTATTTTGGAGTACAAAAACGTTGTAATGGATATAGAAAAATCCTTTAATTCGCAACAACAGCCTTTATGATGAAAGACCAGCTGGGCATACAAAATTTCAAAGCATCCCTTATAAAGTTCTGTAACCGGACAAATAGAACGGTTTCATTCTATCCTCGCTGAAATTAAGATGTTAAAAACTAAACAGGTACACCGAACATTTGAAGAACAGTTCAATTCAGCTGTCTAGGAATATAACTACACAATTCACCCTGTATCAAAAATACAAAACAAAATAACGCCTTAAATAATTTTTGGCAGAAATATAACCACTGATCCAGGAAAATATGACGAAAGCAGATAAGCAACATCGAAAACCTATAATCAAAACAGGCAACAGGCTTAAAAACCATAACGAAAAAGCAATAAGATCAAAGATTATGACCCAGGACAAACAGTTTTTATAAGCAAATCACAAGGCAGGTTCTAAGCTGTACAAGAAAGAACATTAAAAGAAAACAGACAAACATTTGTTATCACAGAAGCAGGAAGATGTGACATACCTGTTCTACTGAAGGCT

compute hash values for spaced seeds in reference (on both strands and for one of the six templates) and perform lookups in hash tables of the reads

Why does MAQ* use **spaced seeds** for mapping?



MAQ uses seed pairs as it allows per default up to 2 **mismatches** between seed and reference. As each 28 mere is represented by 4 non-overlapping seeds, we are guaranteed that we **always have at least 2 seeds** that must result in a perfect match to the reference.

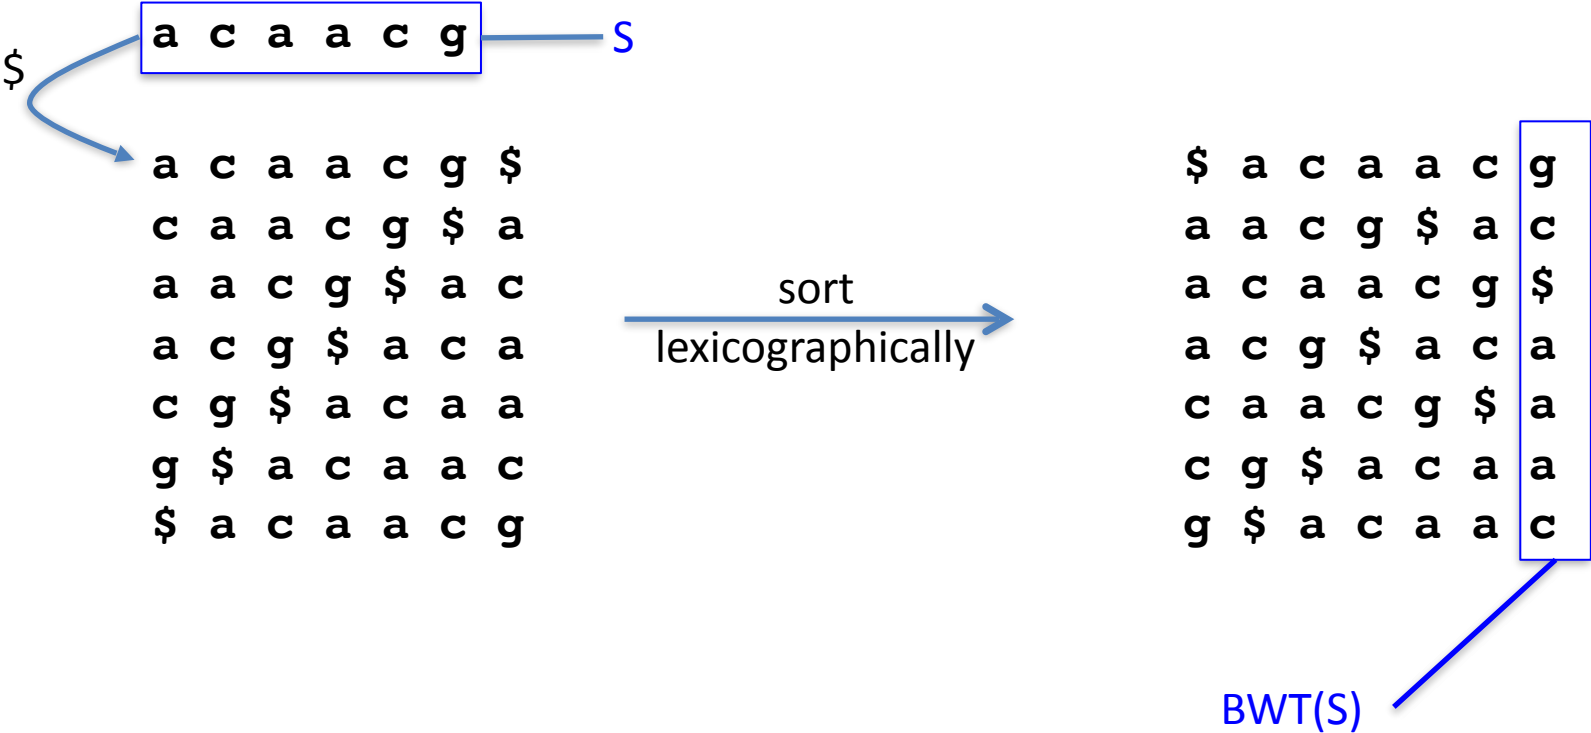
MAQ: An overview

1. At the alignment stage, MAQ first searches for the **ungapped match** with lowest mismatch score, defined as the sum of qualities at mismatching bases.
2. MAQ only considers positions that have **two or fewer mismatches in the first 28 bp** (default parameters; speed-up).
3. Sequences that fail to reach a mismatch score threshold but whose mate pair is mapped are **searched with a gapped alignment algorithm** in the regions defined by the mate pair.
4. To evaluate the reliability of alignments, MAQ assigns each individual alignment a **phred-scaled quality score** (capped at 99), which measures the probability that the true alignment is not the one found by MAQ.
5. MAQ **always reports a single alignment**, and if a read can be aligned equally well to multiple positions, MAQ will **randomly pick** one position and give it a mapping quality zero. Because their mapping score is set to zero, reads that are mapped equally well to multiple positions will not contribute to variant calling.

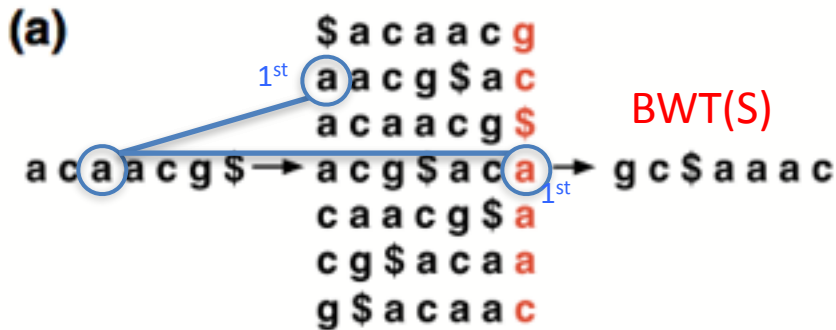
The “new way”: Burrows-Wheeler Transform

- Invented by David Wheeler in 1983 (bell labs), pub. 1994
- Used in data compression (bzip2)
- Used in 2003 on the human genome to define exact word matches (originally for microarray probe design)
- First used for short read alignment by bowtie, now adopted by bwa (maq author) and SOAP2

The Burrows-Wheeler Transform



Burrows Wheeler Transform (BWT)




Generate matrix by

1. Appending a \$ to the end of the string S that should be indexed. \$ should have 2 properties
 1. it must not occur in the string
 2. it should be lexicographically smaller than any character in S
2. generate all cyclic permutations of S
3. sort the resulting matrix lexicographically (the line beginning with the \$ is the first to occur in the matrix).

The matrix has the property of last first (LF) mapping: The *i*th occurrence of character X in the last column corresponds to the same text character as the *i*th occurrence of X in the first column

Burrows Wheeler Transform (BWT)



The matrix has the property of **last first (LF) mapping**: This can be used to reconstruct the original text from BWT(S) using the UNPERMUTE algorithm.

Burrows Wheeler Transform (BWT)



The matrix has the property of last first (LF) mapping: This can be used to search for a text within BWT(S) using the EXACTMATCH algorithm.


Key aspects:

- 1) Matrix is sorted lexicographically. Thus, rows beginning with a given sequence appear **consecutively**.
- 2) EXACTMATCH algorithm calculates the range of matrix rows beginning with **successively longer suffixes** of the query.
- 3) At each step, the size of the **range either shrinks or remains the same**.
- 4) When the algorithm completes, rows beginning with S_0 (the entire query) correspond to **exact occurrences of the query** in the text.

How to cope with mismatches?

Query GGTA

Exact



Empty range:
Abort or **backtrack**.

Ranges of the matrix rows beginning with the suffix observed to that point.

Bowtie conducts a quality-aware, greedy, randomized, depth-first search through the space of possible alignments.



1. The search proceeds similarly to EXACTMATCH
2. If range becomes empty (suffix does not occur in text), the algorithm **backtracks** and selects an already-matched query position and substitutes a different base there. The EXACTMATCH algorithm resumes from this modified position.
3. The algorithm allows only **substitutions that yield a modified suffix that occurs at least once in the text**. If there are multiple candidate substitution positions, then the algorithm greedily selects a position with a **minimal quality** value.
4. Because search is greedy, the **first valid alignment is not necessarily the best** (Number of mismatches and quality). Bowtie has parameters to cope with this (--best or --all (all alignments)).
5. Excessive Backtracking should be avoided. Note, we start from the **low quality** end...

Avoidance of excessive backtracking (assuming 1 mismatch)

Problem: The aligner spends most of its effort fruitlessly backtracking to positions close to the 3' end of the query (error prone).

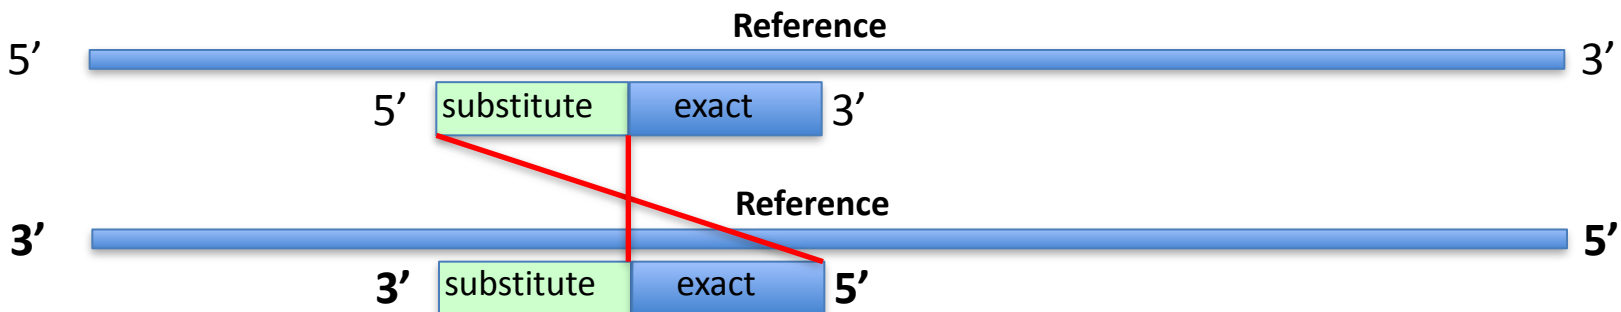
Solution Part1: double indexing (similar to MAQ), using two indices for the genome

- Index 1: BWT of the original genome
- Index 2: BWT of the genome with **reversed character order (not reverse complemented!)**

Solution Part2: The aligner is invoked twice

- **First round:** Index 1 is used, and the aligner is started with original read with the constraint that it must not substitute a position in the query's right half (3' end).
- **Second round:** Index 2 is used, and the aligner is started with the **reversed read**, again with the constraint that it must not substitute a position in the reversed query's right half (originally and still the 5' end).

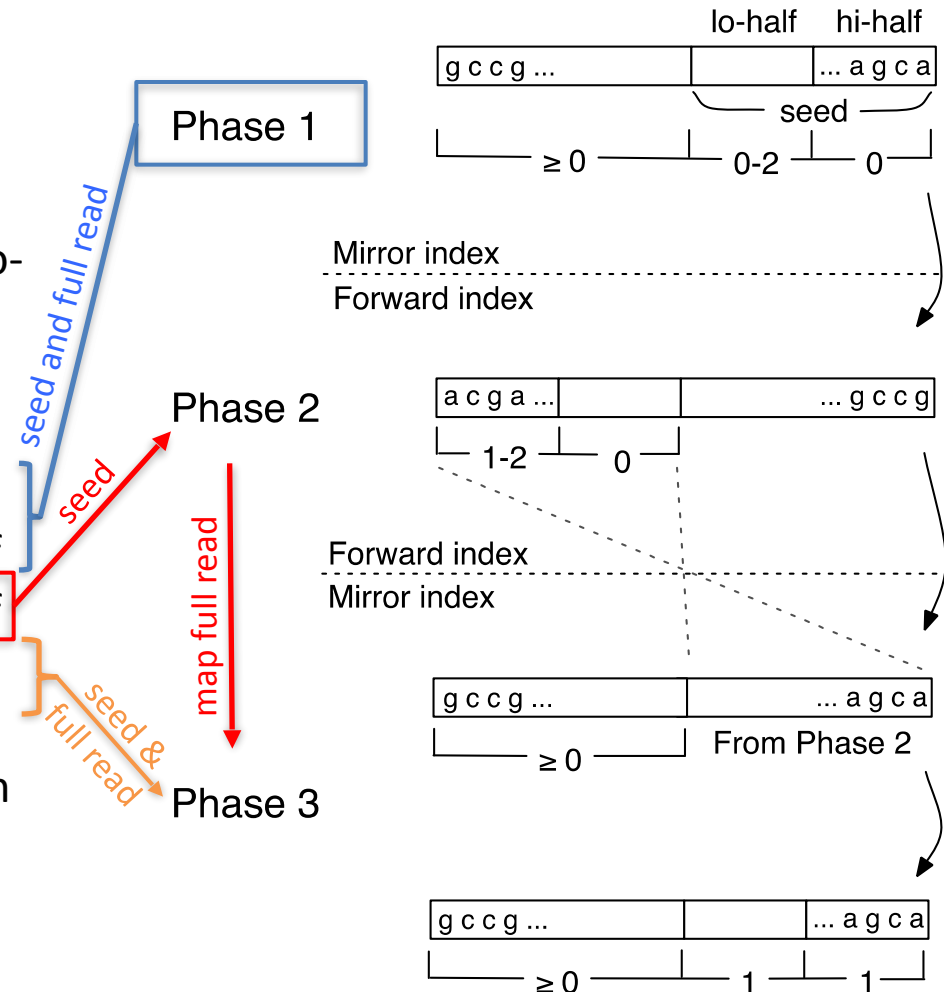
Solution Part3: set a hard upper limit of backtracks to be performed.



The three phases of Bowtie

In the case of 2 (or more) mismatches):

- Bowtie uses the first 28 bp as seed
- The seed is split into a **high quality 5'** half (hi-half) and a **low quality 3'** half (lo-half)
- For up to 2 mismatches we have four scenarios:
 1. no mismatches in seed
 2. 1-2 mismatches only in the lo-half
 3. 1-2 mismatches only in the hi-half
 4. 1 mismatch each in hi- and low-half
- Any number of mismatches can occur in non-seed part (subject to other thresholds).



SAM/BAM format* & SAMtools

Header section

```
@HD VN:1.0
@SQ SN:chr20 AS:HG18 LN:62435964
@RG ID:L1 PU:SC_1_10 LB:SC_1 SM:NA12891
@RG ID:L2 PU:SC_2_12 LB:SC_2 SM:NA12891
```

position of alignment

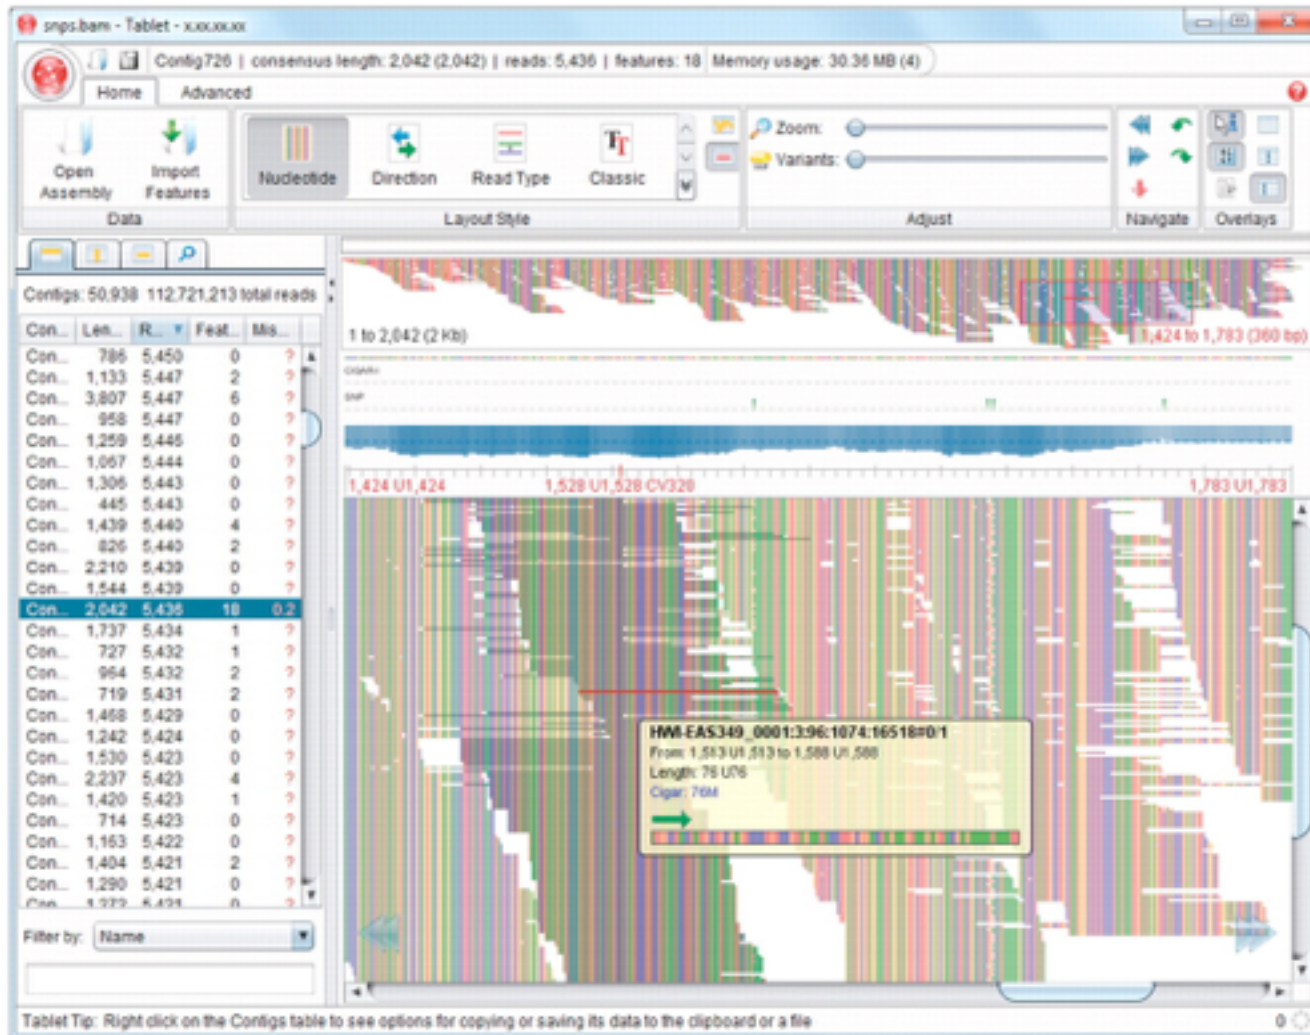
Alignment section

Query Name	Ref sequence	↓	query sequence (same strand as ref)	query quality	
V00-HWI-EAS132:3:38:959:2035#0	147 chr1 28	255	36M = 79 0	GATCTGATGGCAGAAAACCCCTCTCAGTCCGTCGTG	aaX`[\`Y^Y^]ZX``\EV_BBBBBBBBBBBBBBBB NM:i:1
V00-HWI-EAS132:4:99:122:772#0	177 chr1 32	255	36M = 9127 0	AAAGGATCTGATGGCAGAAAACCCCTCTCAGTCCGT	aaaaaa\OwaI_\NL\aa`Xa^\ZUaa[XWT\^XR NM:i:1
V00-HWI-EAS132:4:44:473:970#0	25 chr1 40	255	36M * 0 0	GTCGTGCTGAAGGATCTGATGGCAGAAAACACCTCT	__YaZ`W[aZNUZ[U[_TL[KVVX^QURUTDRVZBB NM:i:2
V00-HWI-EAS132:4:29:113:1934#0	99 chr1 44	255	36M = 107 0	GGTTTTCTGCCATCAGATCCTTTACCACGACAGAC	aaaQaa_``)\`_``^a``a`_``^`_XQ[ZS\XX NM:i:1

- Sequence Alignment/Map & Binary Alignment/Map (Heng Li, et al)
- SAMtools - C or Java
 - convert formats
 - sort/merge alignments
 - generate per-position information
 - call SNPs/indel variants
 - show alignments in a text viewer
 - <http://bioinformatics.oxfordjournals.org/cgi/reprint/btp352v1>

*for details on the format see <http://samtools.sourceforge.net/SAM1.pdf>

Data visualization using Tablet



Mapping: Watch out for the following...

- Up to 2 mismatches in first 28 bases reported by default (Maq, Bowtie, BWA). Maq additionally reports some 3-mismatches.
- If ≥ 1 exact matches exist for a read, Bowtie will report one of them. If the best match is inexact, bowtie will not always return the highest quality alignment (unless run with `--best`).
- Bowtie, BWA, and SOAP2 all randomly place repeat reads.
- Although in theory BWA works with arbitrarily long reads, its performance is degraded on long reads (probably with any BWT aligner). Consider using Bowtie2.
- Non-ACTG bases in the genome are converted to random bases in BWA
- How does the mapper deal with paired end reads? Will it report a mapping success only when both reads map in the correct distance (specify insert size) and orientation?