

EMBL Short Reads Course June 2009:
Managing sequence and annotation data using the
Biostrings and **BSgenome** packages

Patrick Aboyoun
Fred Hutchinson Cancer Research Center
Seattle, WA 98008

6 June 2009

Contents

1 Preliminaries	1
2 Setup	2
3 Basic containers	2
3.1 DNASTring objects	2
3.2 DNASTringSet objects	3
3.3 XStringViews objects	7
4 BSgenome data packages	8
5 String matching	11
5.1 The <code>matchPattern</code> function	11
5.2 The <code>vmatchPattern</code> function	12
5.3 Ambiguities	13
5.4 Masking	13
5.5 Finding the hits of a large set of short motifs	16
6 More on the provenance of the <code>topReads</code> object	19
7 Session Information	20

1 Preliminaries

This lab is designed to teach the basics of `Biostrings` and `BSgenome` data packages. For this lab you need:

- R version 2.9,
- the `Biostrings`, `BSgenome` and `BSgenome.Mmusculus.UCSC.mm9` packages,
- `topReads.rda`: an example data file containing the top 1000 reads for all 8 Solexa lanes of two ChIP-seq experiments. This will be provided to you by the course instructors. Store it on your local file system.

2 Setup

Exercise 1

Start an R session and use the `library` function to load the `BSgenome.Mmusculus.UCSC.mm9` genome package.

```
> library("BSgenome.Mmusculus.UCSC.mm9")
```

Exercise 2

Use the `load` function to load the example dataset into your R session. (You will need to adapt the directory path to where the file is on your system.)

```
> load(file.path("../data", "topReads.rda"))
> ls()

[1] "af"           "chr1"         "dss"          "m"
[5] "m1"           "m2"           "minus_strand" "pattern"
[9] "pdict1"       "pdict2"       "pdss"         "plus_strand"
[13] "r1"           "reg"          "topReads"     "v3"
```

`topReads` is a list of length 2, corresponding to the two experiments. Each element is again a list of length 8, corresponding to the lanes of the Solexa instrument. Each of the elements of that is an `XDataFrame` object with 1000 rows and 2 columns. More on the provenance of the `topReads` object is described in Section 6.

```
> topReads[[1]][[1]]

XDataFrame object with 1000 rows and 2 columns.

> colnames(topReads[[1]][[1]])

[1] "read" "count"

> topReads[[1]][[1]][1,"read"]

A DNASTringSet instance of length 1
width seq
[1] 36 AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

> topReads[[1]][[1]][1,"count"]

[1] 81237
```

3 Basic containers

3.1 DNASTring objects

The `DNASTring` class is the basic container for storing a long nucleotide sequence. Unlike a standard character vector in R that can store multiple strings, a `DNASTring` object can only contain one.

Exercise 3

1. Create a `DNASTring` object `r1` by using the `[[` operator to extract the first read from experiment 2, lane 1.
2. Use the `nchar` and `alphabetFrequency` functions to obtain the number of characters and the frequencies of the different letters in `r1`.

3. Get its reverse complement.

4. Extract substrings with the function `subseq`.

```
> r1 <- topReads[["experiment2"]][["lane1"]][,"read"][[1]]
> nchar(r1)
```

```
[1] 36
```

```
> alphabetFrequency(r1)
```

```
  A  C  G  T  M  R  W  S  Y  K  V  H  D  B  N  -  +
  8  8 10 10  0  0  0  0  0  0  0  0  0  0  0  0  0
```

Note that the *DNAStrng* class can contain characters from the complete set of IUAPC nucleic acid codes (for example, R stands for purine, i.e., A or G).

```
> reverseComplement(r1)
```

```
 36-letter "DNAStrng" instance
seq: TTTCAAGCAGAAGACGGCATACGAGCTCTTCCGATC
```

```
> subseq(r1, start=5, end=15)
```

```
 11-letter "DNAStrng" instance
seq: GGAAGAGCTCG
```

```
> subseq(r1, end=15)
```

```
 15-letter "DNAStrng" instance
seq: GATCGGAAGAGCTCG
```

```
> subseq(r1, start=-5)
```

```
  5-letter "DNAStrng" instance
seq: TGAAA
```

3.2 DNAStrngSet objects

The *DNAStrngSet* class is the basic container for storing multiple nucleotide sequences. As with R vectors, the `length` function returns the number of elements (sequences) in a *DNAStrngSet* object and the `[` operator can be used to subset it. In addition, the element access operator `[[` can be used to extract a single element and return it as a *DNAStrng* object.

Exercise 4

1. Use the *DNAStrngSet* constructor to store the 1000 reads from experiment 2, lane 1 into a *DNAStrngSet* object. Let us call this instance `dss`.
2. Use `length` and `width` on `dss`.
3. Use subsetting operator `[` to remove its second element.
4. Use the `rev` to invert the order of its elements.
5. Use subsetting operator `[[` to extract its first element as a *DNAStrng* object.

6. Use the `DNASet` constructor (i) to remove the last 2 nucleotides of each element, then (ii) to keep only the last 10 nucleotides.
7. Call `alphabetFrequency` on `dss` and on its reverse complement. Try again with setting its argument `collapse=TRUE`.
8. Remove reads with `Ns`, and put the “cleaned up” set of sequences back into `dss`.

```
> dss <- topReads[["experiment2"]][["lane1"]][, "read"]
> length(dss)
```

```
[1] 1000
```

```
> table(width(dss))
```

```
 36
1000
```

```
> dss[-2]
```

```
A DNASet instance of length 999
```

```
width seq
```

```
[1] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTAAAA
[2] 36 AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
[3] 36 ANNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
[4] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTGGAT
[5] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTTGAT
[6] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTATAT
[7] 36 GNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
[8] 36 CNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
[9] 36 TNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
... ..
[991] 36 TGTCCACTGTAGGACGTGGAATATGGCAAGAAAAC
[992] 36 ATTCCTCCGACACATAATAATCAGAACAACAAATG
[993] 36 ATTGATATACTGTTCTACAAATCCCGTTTCCAAC
[994] 36 ANNNNNNNNAAAAANNNANNAAAAAAAAAAAAAAA
[995] 36 ANNNNNNNNNNNNNNNNNNNNNNNAANNANNNNN
[996] 36 CATATTCCAGGTCCTACAGTGTGATTTCTCATTTT
[997] 36 CNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNTN
[998] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTGGAT
[999] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGTTTAGA
```

```
> rev(dss)
```

```
A DNASet instance of length 1000
```

```
width seq
```

```
[1] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGTTTAGA
[2] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTGGAT
[3] 36 CNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNTN
[4] 36 CATATTCCAGGTCCTACAGTGTGATTTCTCATTTT
[5] 36 ANNNNNNNNNNNNNNNNNNNNNNNAANNANNNNN
[6] 36 ANNNNNNNNAAAAANNNANNAAAAAAAAAAAAAAA
[7] 36 ATTGATATACTGTTCTACAAATCCCGTTTCCAAC
```

```

[8] 36 ATTCCTCCCGACACATAATAATCAGAACAACAAATG
[9] 36 TGTCCACTGTAGGACGTGGAATATGGCAAGAAAAC
... ..
[992] 36 CNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
[993] 36 GNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
[994] 36 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTATAT
[995] 36 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTTGAT
[996] 36 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGGAT
[997] 36 ANNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
[998] 36 AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
[999] 36 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTAGAT
[1000] 36 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA

```

```
> dss[[1]]
```

```

36-letter "DNAStrng" instance
seq: GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA

```

```
> DNAStrngSet(dss, end=-3)
```

```

A DNAStrngSet instance of length 1000
width seq
[1] 34 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGA
[2] 34 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTAG
[3] 34 AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
[4] 34 ANNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
[5] 34 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGG
[6] 34 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTTG
[7] 34 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTAT
[8] 34 GNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
[9] 34 CNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
... ..
[992] 34 TGTCCACTGTAGGACGTGGAATATGGCAAGAAA
[993] 34 ATTCCTCCCGACACATAATAATCAGAACAACAAA
[994] 34 ATTGATATACTACTGTTCTACAAATCCCGTTTCCA
[995] 34 ANNNNNNNNNNAAAAANNNNANNAAAAAAAAAAAAA
[996] 34 ANNNNNNNNNNNNNNNNNNNNNNNNNAANNANNNN
[997] 34 CATATTCCAGGTCTACAGTGTGCATTTCTCATT
[998] 34 CNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
[999] 34 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTGG
[1000] 34 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTGA

```

```
> DNAStrngSet(dss, start=-10)
```

```

A DNAStrngSet instance of length 1000
width seq
[1] 10 CTGCTTGAAA
[2] 10 CTGCTTAGAT
[3] 10 AAAAAAAAAA
[4] 10 NNNNNNNNNN
[5] 10 CTGCTTGGAT
[6] 10 CTGCTTTGAT

```

```

[7] 10 CTGCTTATAT
[8] 10 NNNNNNNNNN
[9] 10 NNNNNNNNNN
... ..
[992] 10 CAAGAAAAC
[993] 10 ACAACAAATG
[994] 10 CGTTTCCAAC
[995] 10 AAAAAAAAAA
[996] 10 NNNANNNNNN
[997] 10 TTCTCATTTT
[998] 10 NNNNNNNNTN
[999] 10 CTGCTGGAT
[1000] 10 TCTGTTTAGA

```

```
> head(alphabetFrequency(dss))
```

```

      A C G T M R W S Y K V H D B N - +
[1,]  8 8 10 10 0 0 0 0 0 0 0 0 0 0 0 0 0
[2,]  7 8 10 11 0 0 0 0 0 0 0 0 0 0 0 0 0
[3,] 36 0  0  0 0 0 0 0 0 0 0 0 0 0 0 0 0
[4,]  1 0  0  0 0 0 0 0 0 0 0 0 0 0 35 0 0
[5,]  6 8 11 11 0 0 0 0 0 0 0 0 0 0 0 0 0
[6,]  6 8 10 12 0 0 0 0 0 0 0 0 0 0 0 0 0

```

```
> reverseComplement(dss)
```

```

A DNASTringSet instance of length 1000
width seq
[1] 36 TTCAAGCAGAAGACGGCATAACGAGCTCTCCGATC
[2] 36 ATCTAAGCAGAAGACGGCATAACGAGCTCTCCGATC
[3] 36 TTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTT
[4] 36 NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNT
[5] 36 ATCCAAGCAGAAGACGGCATAACGAGCTCTCCGATC
[6] 36 ATCAAAGCAGAAGACGGCATAACGAGCTCTCCGATC
[7] 36 ATATAAGCAGAAGACGGCATAACGAGCTCTCCGATC
[8] 36 NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNC
[9] 36 NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNG
... ..
[992] 36 AGTTTCTTGCCATATTCCACGTCCTACAGTGGACA
[993] 36 CATTGTTGTTCTGATTATTATGTGTCGGGAGGAAT
[994] 36 GTTGAAAACGGGATTTGTAGAACAGTGTATATCAAT
[995] 36 TTTTTTTTTTTTTTNNNNNNNTTTTTNNNNNNNNNT
[996] 36 NNNNNNTNNNTTNNNNNNNNNNNNNNNNNNNNNNNT
[997] 36 AAAATGAGAAATGCACACTGTAGGACCTGGAATATG
[998] 36 NANNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNG
[999] 36 ATCCAAGCAGAAGGCGGCATAACGAGCTCTCCGATC
[1000] 36 TCTAAACAGAAGACCGGCATAACGAGCTCTCCGATC

```

```
> # Use 'collapse=TRUE' to collapse all the rows
```

```
> alphabetFrequency(dss, collapse=TRUE)
```

```

      A C G T M R W S Y K V H D B
9713 5970 6197 8955 0 0 0 0 0 0 0 0 0

```

```

      N   -   +
5165   0   0

> alphabetFrequency(reverseComplement(dss), collapse=TRUE)

      A   C   G   T   M   R   W   S   Y   K   V   H   D   B
8955 6197 5970 9713   0   0   0   0   0   0   0   0   0   0
      N   -   +
5165   0   0

> # Use [, ] to subset the matrix returned by alphabetFrequency()
> dss <- dss[alphabetFrequency(dss)[, "N"] == 0]

```

3.3 XStringViews objects

An *XStringViews* object contains a set of views on the same sequence, which is called the *subject*; for example, this can be a *DNAStrng* object. Each view is defined by its start and end locations: both are integers such that $\text{start} \leq \text{end}$. The `Views` function can be used to create an *XStringViews* object, given a subject and a set of start and end locations. Like for *DNAStrngSet* objects, `length`, `width`, `[` and `[[` are supported for *XStringViews* objects. Additional methods `subject`, `start`, `end` and `gaps` are also provided.

A typical use case for views is for the subject to be the sequence of a molecule (e.g. a chromosome) and the different views to be certain features of the sequence, such as protein-binding regions, transcribed regions, etc.

Exercise 5

1. Use the `Views` function to create an *XStringViews* object with a *DNAStrng* subject. Make it such that some views are overlapping but also that the set of views do not cover the subject entirely.
2. Try `subject`, `start`, `end` and `gaps` on this *XStringViews* object.
3. Try `alphabetFrequency` on it.
4. Turn it into a *DNAStrngSet* object with the *DNAStrngSet* constructor.

```

> v3 <- Views(dss[[1]], start=c(2, 12, 20), end=c(5, 26, 27))
> subject(v3)

```

```

36-letter "DNAStrng" instance
seq: GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGA AAA

```

```

> start(v3)

```

```

[1]  2 12 20

```

```

> end(v3)

```

```

[1]  5 26 27

```

```

> gaps(v3)

```

```

Views on a 36-letter DNAStrng subject
subject: GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGA AAA
views:

```

```

      start end width
[1]     1   1     1 [G]
[2]     6  11     6 [GAAGAG]
[3]    28  36     9 [TGCTTGA AAA]

```

```
> alphabetFrequency(v3)
      A C G T M R W S Y K V H D B N - +
[1,] 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0
[2,] 1 5 3 6 0 0 0 0 0 0 0 0 0 0 0 0 0 0
[3,] 0 4 1 3 0 0 0 0 0 0 0 0 0 0 0 0 0 0
```

```
> DNASTringSet(v3)
A DNASTringSet instance of length 3
width seq
[1] 4 ATCG
[2] 15 CTCGTATGCCGTCTT
[3] 8 CCGTCTTC
```

4 BSgenome data packages

The name of a BSgenome data package is made of 4 parts separated by a dot, for example BSgenome.Celegans.UCSC.ce2.

1. The 1st part is always the prefix BSgenome.
2. The 2nd part is the (abbreviated) name of the organism.
3. The 3rd part is the name of the organisation who assembled the genome.
4. The 4th part is the release string or number used by this organisation for this assembly of the genome.

All BSgenome data packages contain a single top level object whose name matches the second part of the package name.

Exercise 6

1. Load *BSgenome.Mmusculus.UCSC.mm9* and display its top level object. Note that this does not load any sequence into memory yet.
2. Use *seqlengths* on it to get the lengths of the single sequences (this does not load any sequence either).

```
> Mmusculus
Mouse genome
|
| organism: Mus musculus (Mouse)
| provider: UCSC
| provider version: mm9
| release date: Jul. 2007
| release name: NCBI Build 37
|
| single sequences (see '?seqnames'):
| chr1      chr2      chr3      chr4      chr5
| chr6      chr7      chr8      chr9      chr10
| chr11     chr12     chr13     chr14     chr15
| chr16     chr17     chr18     chr19     chrX
| chrY      chrM      chr1_random chr3_random chr4_random
| chr5_random chr7_random chr8_random chr9_random chr13_random
```



```
| chr16_random chr17_random chrX_random chrY_random chrUn_random
|
| multiple sequences (see '?mseqnames'):
| upstream1000 upstream2000 upstream5000
|
| (use the '$' or '[' operator to access a given sequence)
```

```
> seqlengths(Mmusculus)
```

```
chr1 chr2 chr3 chr4 chr5
197195432 181748087 159599783 155630120 152537259
chr6 chr7 chr8 chr9 chr10
149517037 152524553 131738871 124076172 129993255
chr11 chr12 chr13 chr14 chr15
121843856 121257530 120284312 125194864 103494974
chr16 chr17 chr18 chr19 chrX
98319150 95272651 90772031 61342430 166650296
chrY chrM chr1_random chr3_random chr4_random
15902555 16299 1231697 41899 160594
chr5_random chr7_random chr8_random chr9_random chr13_random
357350 362490 849593 449403 400311
chr16_random chr17_random chrX_random chrY_random chrUn_random
3994 628739 1785075 58682461 5900358
```

Display information about the mitochondrial chromosome.

```
> Mmusculus$chrM
```

```
16299-letter "MaskedDNASTring" instance (# for masking)
seq: GTTAATGTAGCTTAATAACAAAGCAAAGCACT...ATCATACTCTATTACGCAATAAACATTAACAA
masks:
maskedwidth maskedratio active names
1 0 0.00000000 TRUE AGAPS
2 0 0.00000000 TRUE AMB
3 414 0.02540033 FALSE RM
4 0 0.00000000 FALSE TRF
desc
1 assembly gaps (empty)
2 intra-contig ambiguities (empty)
3 RepeatMasker
4 Tandem Repeats Finder [period<=12] (empty)
all masks together:
maskedwidth maskedratio
414 0.02540033
all active masks together:
maskedwidth maskedratio
0 0
```

Some information about the built-in masks is displayed. Let us drop the masks for now by accessing the sequence with

```
> unmasked(Mmusculus$chrM)
```




Figure 1: Nucleotide frequencies.

```
16299-letter "DNAStrng" instance
seq: GTTAATGTAGCTTAATAACAAAGCAAAGCACT...ATCATACTCTATTACGCAATAAACATTAACAA
```

Exercise 7

1. Do the chromosomes contain IUPAC extended letters?
2. Treatment of DNA with bisulfite converts cytosine residues to uracil, but leaves 5-methylcytosine residues unaffected. The first reported method [1] of methylation analysis using bisulfite-treated DNA used PCR and standard dideoxynucleotide DNA sequencing to directly determine the nucleotides resistant to bisulfite conversion. All sites of unmethylated cytosines were displayed as thymines in the resulting amplified sequence of the sense strand, and as adenines in the amplified antisense strand. Assuming that no part of it is methylated, use the `chartr` function to simulate a bisulfite transformation of a 5 kb segment of chromosome 1.

Apply `alphabetFrequency` to each unmasked chromosome:

```
> af <- sapply(seqnames(Mmusculus),
+             function(name)
+             alphabetFrequency(unmasked(Mmusculus[[name]]), baseOnly=TRUE))
```

and plot the result as a barplot (see Fig. 1).

```
> barplot(t(af), beside=TRUE, col=rainbow(ncol(af)))
```

Bisulfite transformation of the plus strand:

```
> reg = 10000000+(1:5000)
> plus_strand <- chartr("C", "T", unmasked(Mmusculus$chr1)[reg])
> alphabetFrequency(plus_strand, baseOnly=TRUE)
```

A	C	G	T	other
1495	0	1061	2444	0

5.2 The vmatchPattern function

This function finds all the matches of a given pattern in a set of reference sequences (the “v” stands for *vectorized*).

Exercise 9

1. Load the `upstream5000` object from `Mmusculus` and find all the matches of a short pattern in it.
2. The value returned by `vmatchPattern` is an `MIndex` object containing the match coordinates for each reference sequence. You can use the `startIndex` and `endIndex` accessors on it to extract the match starting and ending positions as lists (one list element per reference sequence). `[[` extracts the matches of a given reference sequence as an `MIndex` object. `countIndex` extracts the match counts as an integer vector (one element per reference sequence).

```
> Mmusculus$upstream5000
```

```
A DNASTringSet instance of length 18429
      width seq
[1] 5000 AGGAAGAACATATTCTC...GAACGCGGGGCTTTCTA NM_028778_up_5000...
[2] 5000 ATCCCAAAAGTCCCCCA...TCTTCAGCTGGAGCTGG NM_027671_up_5000...
[3] 5000 TTCTTTACTTAGAAAAGT...ACTTGGATAAGGCGCAA NM_175642_up_5000...
[4] 5000 TGGGTCAAGCATACAAA...CTCCCGCCACTGGGAGA NM_008922_up_5000...
[5] 5000 GTAGCCCAAGTGCTCAG...CCATCCTGGGGCACAAG NM_175370_up_5000...
[6] 5000 ATGAAACCACTATGATA...CGCGAGCCTGACGTTGC NM_178884_up_5000...
[7] 5000 TTGTGTGCATCATTTCA...CTGCTAACTTCTGCCTT NM_009126_up_5000...
[8] 5000 ATTAACCTGATCCTGAT...GCCACACACAGGCTTCT NM_198680_up_5000...
[9] 5000 AGCAGAGAGACTCTTTC...GCTTTTCTCTTCCGCCA NM_199021_up_5000...
...
[18421] 5000 TTAAGAACTTTACGCT...TTTTTTTTTTTGCCATT NM_001037748_up_5...
[18422] 5000 GCCATTCCAAAAAAGTT...GGACTGAAGGTGGAGG NM_011667_up_5000...
[18423] 5000 TGCATTAGGCACACATA...TTCAAGGTGAGTTCACT NM_001017393_up_5...
[18424] 5000 AAGAGAAATAATTGATC...TTTTTTTTTTTGCCATT NM_001037748_up_5...
[18425] 5000 GTGGGTGTTAGAAATTG...GCGCATCTATTCCACTT NM_001025241_up_5...
[18426] 5000 ACTATTGATCCTTAGGC...ACTTAGAGACTAGAA NM_009220_up_5000...
[18427] 5000 TTGATCCTCACTAAAAT...TTTTTTTTTTTGCCATT NM_001037748_up_5...
[18428] 5000 TGATCCTCACTAAAAT...TTTTTTTTTTTGCCATT NM_001037748_up_5...
[18429] 5000 CCATGTGGGTGTTAGAA...GCGCATCTATTCCACTT NM_001025241_up_5...
```

```
> m <- vmatchPattern(pattern, Mmusculus$upstream5000)
```

To get the indices of the references sequences with hits:

```
> which(countIndex(m) != 0)
```

```
[1] 2956 7540 10701 11387
```

To get the hits in reference sequence 2956:

```
> m[[2956]]
```

IRanges object:

```
      start end width
[1] 3682 3691    10
```


Exercise 11

What percentage of Mouse chromosome 1 is made of assembly gaps?

```
> maskedratio(masks(Mmusculus$chr1) ["AGAPS"])
```

```
[1] 0.02899639
```

Exercise 12

Check the alphabet frequency of Mouse chromosome 1 when only the AGAPS mask is active, when only the AGAPS and AMB masks are active. Compare with unmasked chromosome 1.

Mmusculus\$chr1 is an immutable object, so before we can turn its masks on or off, we need to copy it to another variable (note that the chromosome sequence itself is not copied during this operation, so it does not result in the use of a substantial amount of additional memory):

```
> chr1 <- Mmusculus$chr1
> active(masks(chr1)) <- FALSE
> active(masks(chr1)) ["AGAPS"] <- TRUE
> chr1
```

```
197195432-letter "MaskedDNAString" instance (# for masking)
seq: #####...AGAATTTGGTATTAACTTAAACTGGAATTC
masks:
```

	maskedwidth	maskedratio	active	names
1	5717956	2.899639e-02	TRUE	AGAPS
2	47	2.383422e-07	FALSE	AMB
3	84650265	4.292709e-01	FALSE	RM
4	4014755	2.035927e-02	FALSE	TRF

	desc
1	assembly gaps
2	intra-contig ambiguities
3	RepeatMasker
4	Tandem Repeats Finder [period<=12]

all masks together:

maskedwidth	maskedratio
90481616	0.4588424

all active masks together:

maskedwidth	maskedratio
5717956	0.02899639

```
> alphabetFrequency(chr1, baseOnly=TRUE)
```

A	C	G	T	other
56406566	39397656	39371416	56301791	47

```
> active(masks(chr1)) ["AMB"] <- TRUE
```

```
> alphabetFrequency(chr1, baseOnly=TRUE)
```

A	C	G	T	other
56406566	39397656	39371416	56301791	0

```
> alphabetFrequency(unmasked(chr1), baseOnly=TRUE)
```

A	C	G	T	other
56406566	39397656	39371416	56301791	5718003


```

+             tables(readXStringColumns(baseCallPath(sp[[i]]),
+                                     pattern = patSeq[[j]],
+                                     colClasses =
+                                     c(rep(list(NULL), 4),
+                                       list("DNAStrng")))[[1]],
+                   n = n)[["top"]]
+             names(x) <- chartr("-", "N", names(x))
+             cat("done.\n")
+             XDataFrame(read = DNAStrngSet(names(x)),
+                       count = unname(x))
+         })
+     })
> names(topReads) <- names(sp)
> for (i in seq_len(length(sp))) {
+     names(topReads[[i]]) <- names(patSeq)
+ }

```

You could adapt this for use with your own data.

7 Session Information

```
> toLatex(sessionInfo())
```

- R version 2.9.0 (2009-04-17), x86_64-unknown-linux-gnu
- Locale: LC_CTYPE=C;LC_NUMERIC=C;LC_TIME=C;LC_COLLATE=C;LC_MONETARY=C;LC_MESSAGES=it_IT.UTF-8;LC_PAPER=it_IT.UTF-8;LC_NAME=C;LC_ADDRESS=C;LC_TELEPHONE=C;LC_MEASUREMENT=it_IT.UTF-8;LC_IDENTIFICATION=C
- Base packages: base, datasets, grDevices, graphics, methods, stats, utils
- Other packages: BSgenome 1.12.2, BSgenome.Mmusculus.UCSC.mm9 1.3.13, Biostrings 2.12.5, IRanges 1.2.2, fortunes 1.3-6
- Loaded via a namespace (and not attached): Biobase 2.4.1

References

- [1] Frommer M, McDonald LE, Millar DS, Collis CM, Watt F, Grigg GW, Molloy PL, Paul CL. A genomic sequencing protocol that yields a positive display of 5-methylcytosine residues in individual DNA strands. *Proc Natl Acad Sci U S A* 89:1827–1831 (1992)