

# Supplementary material

This supplementary material describes all the necessary steps to repeat the experiments described in the main paper. The package with all source codes can be downloaded from: <http://sun.aei.polsl.pl/mugi/>. Additional tables and figures are also included.

## 1 Multple Genomes Index description

Multiple Genome Index is a compressed index of collections of genomes of the same species. It allows to ask for both exact and approximate queries.

Here we describe the usage of the tools to build the index (**MuGI\_build**), perform the search (**MuGI\_search**), and parse the result to a readable form (**MuGI\_parse**). We also provide a code generating example queries (**MuGI\_exGen**). All these tools are implemented in C/C++.

### 1.1 MuGI\_build

The program builds the index with set sparsity and length of  $k$ -mers, based on the input reference sequence and Variant Call Format (VCF) file describing the differences between individual genomes and reference genome.

Usage:

- **MuGI\_build** `<input.fasta>` `<input.vcf>` `<k>` `<sp>` `<pl>`

Parameters:

- **input.fasta** — name of the input FASTA file with the reference
- **input.vcf** — name of the input VCF file
- **k** — desired length of  $k$ -mers in the index
- **sp** — desired sparsity of the index
- **pl** — ploidity of the individuals in the collection (only 1 and 2 supported)

The output files:

- **index-<k>-<sp>** — built index
- **index-dict** — dictionary file

The following are required to build the sources:

- Intel Threading Building Blocks library (<https://www.threadingbuildingblocks.org>)
- Compiler with OpenMP support (<http://openmp.org>)

If the ploidity is set to 2, any haploid genotype call is treated as diploid. The diploid genotypes should be phased. The first word in successive chromosome headers should match successive chromosomes described in the VCF file (its first column). The positions of the variants in the VCF file should be within the current chromosome (first column). During the process they are adjusted to present the position within the whole reference based on the size of chromosomes sequences. The **index-dict** file is created along with the index. It contains data about sizes of all chromosomes and names of all individuals. It is used to parse the search results to a readable form.

For testing we implemented the possibility to build few indexes with the same k-mer length and different sparsities (from 1 up to  $< sp >$ ) at the same run. To use this mode, `CREATE_ALL_INDEXES` must be defined in the `defs.h` file (see line 14 of `defs.h`).

## 1.2 MuGI\_search

The program performs the search of the input queries in the compressed collection (input index file) and reports to the stdout statistics about queries execution times. Asmlib library (<http://www.agner.org/optimize/asmlib.zip>) is required to build the sources.

Usage:

- `MuGI_search <index-file> <queries-file> {<maxError>}`

Parameters:

- `index-file` — path to file with the index
- `queries-file` — path to file with the sequences to search (FASTA or FASTQ or plain sequences - one in each line)
- `maxError` — optional, maximum number of mismatches in found sequence (default 0, exact search)

The output file:

- `result.out` — binary file with the result

## 1.3 MuGI\_parse

The program parses the binary file with result (output of `MuGI_search`) and creates a text file with readable data.

For each queries, at first its name (for FASTQ and FASTA) or ordinal number (for simple list) is written. Then, all found positions are reported (chromosome number and position within it) along with list of all individuals in which the match was found. In a "hex" mode (0), the list is represented by a bit vector written in a hexadecimal form, where each  $j$ th bit set corresponds to match found in  $j$ th haploid genome. In the "full" mode (1) the list consist of names of all individuals with the match. For diploid individuals, to distinguish between two haploid sequences of each, a suffix "-1" or "-2" is added to the name. In the bit vector representation, each diploid individual is represented by two consecutive bits.

Usage:

- `MuGI_parse <dict> <file_to_parse> <output_file> <mode>`

Parameters:

- `dict` — dictionary file of the index
- `file_to_parse` — binary file with the result
- `output_file` — name of the output file
- `mode` — optional, mode of the output: 0 (hexadecimal individuals representation) or 1 (full individual names) default: 0

## 1.4 MuGI\_genEx

The program generates a file with a set number of queries. The file can be in FASTQ or FASTA format, or be a text file with list of sequences (three acceptable formats of queries for MuGI\_search). It uses the created index file, to take excerpts from the collection of genomes. As only Reference (REF), Variant Database (VD) and Bit Vectors (BV) are used, the  $k$  and  $sp$  parameters of the index building process does not matter. The final queries are made as described in the main paper.

Usage:

- `MuGI_genEx <index-file> <out-file> <no_reads> {<mode>}`

Parameters:

- `index-file` — path to file with the index
- `out-file` — output file with example reads
- `no_reads` — number of reads to generate
- `mode` — optional, output format: "fastq" or "fasta" or "list" (default: "list")

## 2 Data source

The data used in experiments are from the Phase 1 of the 1000 Genomes Project, which contains data about 1092 human individuals.

The human reference sequence used can be found at the NCBI's anonymous FTP server ([ftp://ftp-trace.ncbi.nih.gov/1000genomes/ftp/technical/reference/human\\_g1k.v37.fasta.gz](ftp://ftp-trace.ncbi.nih.gov/1000genomes/ftp/technical/reference/human_g1k.v37.fasta.gz)). We removed all non-chromosomal supercontigs, leaving only 1–22, X, and Y chromosomes.

The VCF (Variant Call Format) files can be downloaded from the NCBI's anonymous FTP server. It can be either EBI ([ftp://ftp.1000genomes.ebi.ac.uk/vol1/ftp/phase1/analysis\\_results/integrated\\_call\\_sets/](ftp://ftp.1000genomes.ebi.ac.uk/vol1/ftp/phase1/analysis_results/integrated_call_sets/)) or NCBI ([ftp://ftp.ncbi.nih.gov/1000genomes/ftp/phase1/analysis\\_results/integrated\\_call\\_sets/](ftp://ftp.ncbi.nih.gov/1000genomes/ftp/phase1/analysis_results/integrated_call_sets/)) FTP site. Note that the size of the complete set is about 1.2 TB. The complete list of compressed VCF files used:

```
ALL.chr1.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr2.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr3.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr4.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr5.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr6.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr7.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr8.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr9.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr10.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr11.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr12.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr13.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr14.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr15.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr16.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr17.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
ALL.chr18.integrated_phase1_v3.20101123.snps_indels_svsvs.genotypes.vcf.gz
```

```

ALL.chr19.integrated_phase1_v3.20101123.snps_indels_svgs.genotypes.vcf.gz
ALL.chr20.integrated_phase1_v3.20101123.snps_indels_svgs.genotypes.vcf.gz
ALL.chr21.integrated_phase1_v3.20101123.snps_indels_svgs.genotypes.vcf.gz
ALL.chr22.integrated_phase1_v3.20101123.snps_indels_svgs.genotypes.vcf.gz
ALL.chrX.integrated_phase1_v3.20101123.snps_indels_svgs.genotypes.vcf.gz
ALL.chrY.phase1_samtools_si.20101123.snps_low_coverage.genotypes.vcf.gz

```

Because of the two pseudoautosomal regions (PAR1 and PAR2) shared between X and Y chromosomes in male individuals, the diploid genotype calls for variants in the pseudoautosomal regions in X and Y chromosomes are stored in the available VCF file for X chromosome, while haploid genotype calls for non-pseudoautosomal regions (nonPAR, between PAR1 and PAR2) of X and Y chromosomes are stored in the corresponding VCF files. For Y chromosome there are information for one additional individual (not present in other VCF files). To have consistent data with one diploid representation for the whole reference, we removed the extra individual data and artificially extended the list of genomes described in the VCF file of chromosome Y to be the same as for other chromosomes, that is, to include female individuals. The alleles values were set to zero (reference allele) in the genotype fields. The tool building the index with the diploidity parameter is set to 2, treats all haploid calls as diploid—occurring variant is introduced in two generated sequences. This modification of the dataset means there can be matches reported in non existing regions, i.e, chromosome Y of a female individual or two chromosomes X of a male individual, but these false results are easy to identify and correct in the post processing process.

To build an index for a single chromosome of all individuals in the collection, the reference sequences of single assembled chromosomes may be useful. They can be found at the GenBank and downloaded from the NCBI's anonymous FTP server:

```

ftp://ftp.ncbi.nlm.nih.gov/genbank/genomes/Eukaryotes/vertebrates_mammals/Homo_sapiens/GRCh37/Primary_Assembly/assembled_chromosomes/FASTA/

```

## 3 Experiments

Here we describe all necessary steps to build an index and perform experiments on single human chromosome and on the whole human genome.

### 3.1 Single chromosome

To build the index for a single autosomal chromosome of all individuals from the collection, it is enough to download the related reference sequence and VCF file. Only FASTA file requires slight modification, so its first word matches **CHROM** field in the VCF file. The script below shows the steps for the 22 chromosome (**wget** and **gzip** tools required). It downloads and decompresses the data, builds the index with sparsity 3 and  $k$ -mer length set to 25, generates 100K queries (in FASTA format), performs exact search and parse the resulting file to store the complete report (positions and names of individuals) in the **full-results.txt** file

---

```

wget ftp://ftp.1000genomes.ebi.ac.uk/vol1/ftp/phase1/analysis_results/
  integrated_call_sets/ALL.chr22.integrated_phase1_v3.20101123.
  snps_indels_svgs.genotypes.vcf.gz
gzip -d ALL.chr22.integrated_phase1_v3.20101123.snps_indels_svgs.genotypes.vcf.gz
mv ALL.chr22.integrated_phase1_v3.20101123.snps_indels_svgs.genotypes.vcf > chr22.vcf
wget ftp://ftp.ncbi.nlm.nih.gov/genbank/genomes/Eukaryotes/vertebrates_mammals/
  Homo_sapiens/GRCh37/Primary_Assembly/assembled_chromosomes/FASTA/chr22.fa.gz
gzip -d chr22.fa.gz

```

```
sed 's/>/>22 /g' chr22.fa > chr22.fasta
./MuGI_build chr22.fasta chr22.vcf 25 3 2
./MuGI_genEx index-25-3 queries.fa 100000 fasta
./MuGI_search index-25-3 queries.fa 0
./MuGI_parse results.out full-results.txt 1
```

---

## 3.2 Whole genome

Building the index for the whole genome requires pre-processing of the available data (reference sequence and the 24 VCF files) as described in Section 2. The detailed steps and additional guidelines (for the unix-like environment):

### 1. Pre-processing the reference sequence

We cut from the downloaded and decompressed reference file anything after the Y chromosome (starting from the mitochondria). The following example script downloads the available reference and creates **human.fasta** file with the desired reference (**wget**, **gzip** and **awk** tools required):

---

```
wget ftp://ftp-trace.ncbi.nih.gov/1000genomes/ftp/technical/reference/
    human_g1k_v37.fasta.gz
gzip -d human_g1k_v37.fasta.gz
awk 'BEGIN{c=1} />MT/ {c=0} // {if(c==1){print $0}}' human_g1k_v37.fasta > human.fasta
```

---

### 2. Pre-processing the Y chromosome

For consistency of the data the VCF file for the Y chromosome were pre-processed as described in Section 2. We are providing the C++ example code **preProcessY.cpp** together with MuGI distribution, that can be used for this step. The arguments of the program are: input Y chromosome VCF file and output name. The following example script downloads and pre-processes the VCF file and creates its modified version: **chrY.vcf** (gcc compiler used in the example, **wget** and **gzip** tools required):

---

```
wget ftp://ftp.1000genomes.ebi.ac.uk/vol1/ftp/phase1/analysis_results/
    integrated_call_sets/ALL.chrY.phase1_samtools_si.20101123.snps.
    low_coverage.genotypes.vcf.gz
gzip -d ALL.chrY.phase1_samtools_si.20101123.snps.low_coverage.genotypes.vcf.gz
g++ -O3 preProcessY/preProcessY.cpp -o prepY
./prepY ALL.chrY.phase1_samtools_si.20101123.snps.low_coverage.genotypes.vcf chrY.vcf
```

---

### 3. Merging the VCF files

We merged all 24 VCF files downloaded from the data source (section 2), including pre-processed VCF file for Y chromosome (previous step) to one VCF file with single header (taken from the VCF file of 1 chromosome, other headers were discarded). The example script performing the task (assuming all VCF files are already decompressed and the modified VCF file name is **chrY.vcf**) and creating single **human.vcf** file:

---

```
cp ALL.chr1.integrated_phase1_v3.20101123.snps_indels_svs.genotypes.vcf human.vcf
tail -n+31 ALL.chr2.integrated_phase1_v3.20101123.snps_indels_svs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr3.integrated_phase1_v3.20101123.snps_indels_svs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr4.integrated_phase1_v3.20101123.snps_indels_svs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr5.integrated_phase1_v3.20101123.snps_indels_svs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr6.integrated_phase1_v3.20101123.snps_indels_svs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr7.integrated_phase1_v3.20101123.snps_indels_svs.genotypes.vcf >> human.vcf
```

---

```

tail -n+31 ALL.chr8.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr9.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr10.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr11.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr12.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr13.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr14.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr15.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr16.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr17.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr18.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr19.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr20.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr21.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+31 ALL.chr22.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+32 ALL.chrX.integrated_phase1_v3.20101123.snps_indels_svcs.genotypes.vcf >> human.vcf
tail -n+27 ALL.chrY.vcf >> human.vcf

```

---

In our experiments we merged the minimal versions of the VCF files (VCFmin files, that can be obtained from the VCF files as in [Deorowicz *et al.*(2013)]), what speeds up the merging and reduces the required disk space. The resulting **human.vcf** file is a valid input for the MuGI\_build program.

#### 4. Building the index

To build the index with  $k$ -mer length 30 and sparsity 4:

```
./MuGI_bulid human.fasta human.vcf 30 4 2
```

---

The resulting index name is **index-30-4**

#### 5. Generating the sample data

To generate 100K sample queries for the collection, stored as a simple list:

```
./MuGI_genEx index-30-4 queries.txt 100000
```

---

#### 6. Searching

Example search for generated sequences in the compressed collection, assuming maximum 2 mismatches:

```
./MuGI_search index-30-4 queries.txt 2
```

---

The result is stored in the binary file **result.out**

#### 7. Parsing the result

To see the results in a readable format (with hexadecimal representation of occurrence or not of match in all individuals):

```
./MuGI_parse index-dict result.out parsed_results.txt 0
```

---

## References

[Deorowicz *et al.*(2013)] Deorowicz, S. *et al.* (2013) Genome compression: a novel approach for large collections. *Bioinformatics*, **29**(20), 2572–2578.

## 4 Additional tables

Table 1: Query times for  $k = 40$  and *sparsity* = 3 (size 13.4 GB). All times are expressed in  $\mu s$

Percentile	Max. allowed mismatches					
	0	1	2	3	4	5
10%	12.4	27.7	44.0	62.1	81.1	100.0
25%	15.4	32.2	50.3	69.9	90.7	111.8
50%	18.7	38.3	58.8	81.4	105.8	131.4
75%	23.8	48.5	73.5	103.7	141.9	199.5
90%	35.9	76.5	116.0	188.0	707.0	3,747.5
95%	57.0	112.7	182.9	718.1	4,972.8	17,619.4
average	39.4	85.5	136.1	372.5	1,334.0	4,021.0

Table 2: Query times for various variants of the index. All times are expressed in  $\mu s$ . For sparsities larger than 8, the times for large values of mismatches (4 and 5) are significant and we do not recommend to use such setting

$k$	$sparsity$	size [GB]	Max. allowed mismatches					
			0	1	2	3	4	5
25	1	24.7	214.2	450.8	699.5	971.5	1,438.3	2,976.8
25	3	11.8	225.0	481.2	751.6	1,024.3	1,599.9	3,647.1
25	4	10.2	229.8	493.1	754.0	1,050.9	1,676.4	4,004.2
25	8	7.7	243.1	528.3	814.6	1,158.5	2,341.4	6,790.4
25	12	6.9	257.2	558.3	868.0	1,337.8		
25	16	6.5	268.8	588.6	916.6	1,787.9		
30	1	26.3	85.4	193.4	303.0	456.0	1,036.4	3,004.6
30	3	12.3	95.7	220.8	340.6	520.4	1,258.0	3,716.2
30	4	10.6	100.4	227.8	351.5	544.0	1,376.5	4,104.5
30	8	7.9	121.8	267.0	414.5	713.6	2,215.2	6,994.5
30	12	7.1	134.0	291.4	456.4	959.4		
30	16	6.6	149.2	319.3	506.8	1,490.4		
35	1	27.9	41.4	98.0	152.0	301.8	1,033.4	3,114.6
35	3	12.9	53.6	121.2	193.0	380.4	1,280.8	3,861.2
35	4	11.0	58.6	130.2	206.3	419.2	1,411.7	4,277.6
35	8	8.2	77.2	166.1	260.3	608.3	2,224.4	7,120.4
35	12	7.2	93.3	196.2	314.6	905.3		
35	16	6.8	107.0	222.2	382.4	1,506.4		
40	1	29.6	28.8	65.2	102.3	291.0	1,109.9	3,348.8
40	3	13.4	39.4	85.5	136.1	372.5	1,334.0	4,021.0
40	4	11.4	43.4	94.4	151.4	412.2	1,471.1	4,461.1
40	8	8.4	61.0	128.9	210.3	615.4	2,297.9	7,350.3
40	12	7.4	76.3	160.0	271.8	917.0		
40	16	6.9	90.4	184.4	344.3	1,514.1		
45	2	18.3	25.9	56.2	97.9	329.6	1,207.0	3,687.2
45	3	14.0	31.3	67.7	116.5	375.5	1,353.3	4,115.0
45	4	11.8	36.3	77.2	132.6	421.2	1,490.9	4,525.3
45	8	8.6	54.2	112.4	196.2	625.8	2,394.4	7,523.9
45	12	7.5	70.4	142.9	262.5	942.0		
45	16	7.0	82.1	168.4	342.3	1,531.9		
GEM mapper		5.0	14.3	26.6	40.4	71.9	126.7	262.7



Table 3: Sizes of the index components. All sizes in MBs

$k$	$sparsity$	REF	VD	BV	$kMA^0$	$kMA^1$	$kMA^2$	$kMA^3$	Total
25	1	1,548	698	2,704	11,502	8,021	84	123	24,680
25	3	1,548	698	2,704	3,879	2,731	84	123	11,767
25	4	1,548	698	2,704	2,926	2,070	84	123	10,153
25	8	1,548	698	2,704	1,496	1,078	84	123	7,732
25	12	1,548	698	2,704	1,020	747	84	123	6,925
25	16	1,548	698	2,704	782	582	84	123	6,521
30	1	1,548	698	2,704	11,502	9,634	85	137	26,307
30	3	1,548	698	2,704	3,879	3,270	85	137	12,320
30	4	1,548	698	2,704	2,926	2,474	85	137	10,571
30	8	1,548	698	2,704	1,496	1,281	85	137	7,948
30	12	1,548	698	2,704	1,020	883	85	137	7,074
30	16	1,548	698	2,704	782	684	85	137	6,637
35	1	1,548	698	2,704	11,502	11,254	85	151	27,942
35	3	1,548	698	2,704	3,879	3,810	85	151	12,875
35	4	1,548	698	2,704	2,926	2,880	85	151	10,992
35	8	1,548	698	2,704	1,496	1,484	85	151	8,167
35	12	1,548	698	2,704	1,020	1,019	85	151	7,225
35	16	1,548	698	2,704	782	786	85	151	6,754
40	1	1,548	698	2,704	11,502	12,881	86	166	29,584
40	3	1,548	698	2,704	3,879	4,354	86	166	13,434
40	4	1,548	698	2,704	2,926	3,288	86	166	11,415
40	8	1,548	698	2,704	1,496	1,689	86	166	8,387
40	12	1,548	698	2,704	1,020	1,156	86	166	7,377
40	16	1,548	698	2,704	782	889	86	166	6,872
45	1	1,548	698	2,704	11,502	14,515	86	181	31,234
45	2	1,548	698	2,704	5,784	7,303	86	181	18,305
45	4	1,548	698	2,704	2,926	3,697	86	181	11,840
45	8	1,548	698	2,704	1,496	1,894	86	181	8,608
45	12	1,548	698	2,704	1,020	1,293	86	181	7,531
45	16	1,548	698	2,704	782	993	86	181	6,992

## 5 Additional figures




Figure 1: Times for exact queries for various  $k$  and *sparsity*. The rightmost points for each  $k$  is for *sparsity* = 1 except for  $k = 45$ , where the rightmost point is for *sparsity* = 2 (due to limitation of our test machine)




Figure 2: Times for up to 1 mismatch queries for various  $k$  and *sparsity*. The rightmost points for each  $k$  is for *sparsity* = 1 except for  $k = 45$ , where the rightmost point is for *sparsity* = 2 (due to limitation of our test machine)




Figure 3: Times for up to 2 mismatches queries for various  $k$  and *sparsity*. The rightmost points for each  $k$  is for *sparsity* = 1 except for  $k = 45$ , where the rightmost point is for *sparsity* = 2 (due to limitation of our test machine)




Figure 4: Times for up to 3 mismatches queries for various  $k$  and *sparsity*. The rightmost points for each  $k$  is for *sparsity* = 1 except for  $k = 45$ , where the rightmost point is for *sparsity* = 2 (due to limitation of our test machine)




Figure 5: Times for up to 4 mismatches queries for various  $k$  and *sparsity*. The rightmost points for each  $k$  is for *sparsity* = 1 except for  $k = 45$ , where the rightmost point is for *sparsity* = 2 (due to limitation of our test machine)




Figure 6: Times for up to 5 mismatches queries for various  $k$  and *sparsity*. The rightmost points for each  $k$  is for *sparsity* = 1 except for  $k = 45$ , where the rightmost point is for *sparsity* = 2 (due to limitation of our test machine)

$k = 40, sparsity = 3$




Figure 7: Percentile times for various types of queries (from 0 to 5 mismatches)



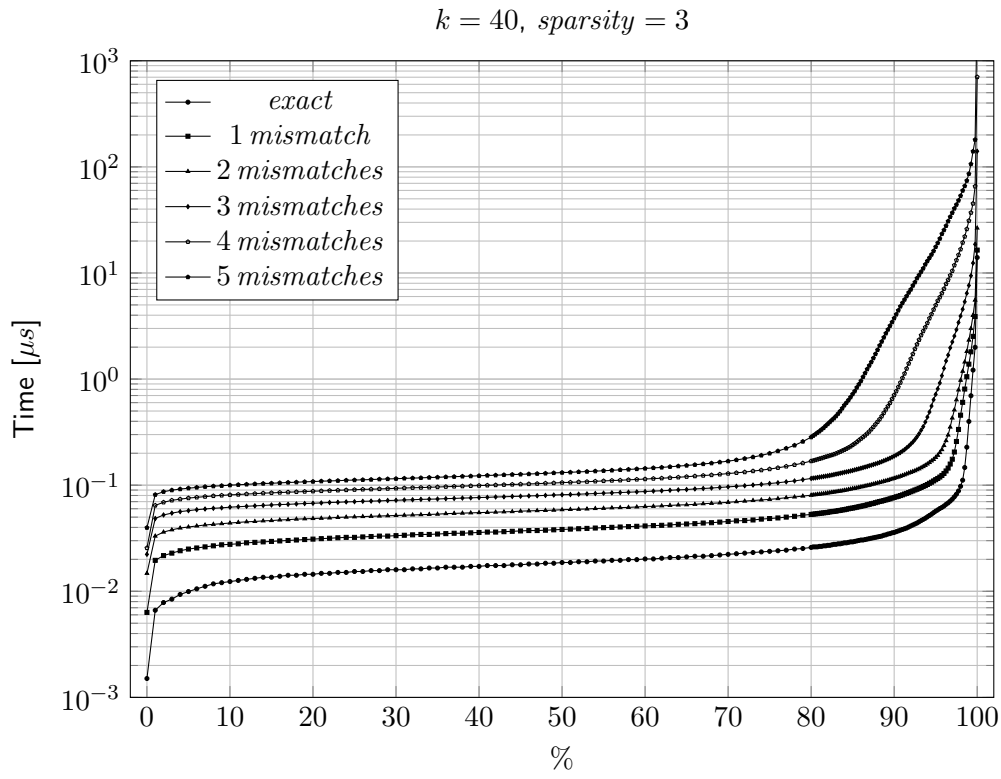


Figure 8: Percentile times for various types of queries (from 0 to 5 mismatches)




Figure 9: Percentile times for various types of queries (from 0 to 5 mismatches)




Figure 10: Percentile times for various types of queries (from 0 to 5 mismatches)




Figure 11: Percentile times for various types of queries (from 0 to 5 mismatches)




Figure 12: Percentile times for various types of queries (from 0 to 5 mismatches)

$k = 35$ , *exact*




Figure 13: Percentile times for various types of queries (from 0 to 5 mismatches)

$k = 35, errors = 3$




Figure 14: Percentile times for various types of queries (from 0 to 5 mismatches)




Figure 15: Percentile times for various types of queries (from 0 to 5 mismatches)