arXiv:1607.06618v1 [cs.DS] 22 Jul 2016

Gerbil: A Fast and Memory-Efficient k-mer Counter with GPU-Support*

Marius Erbert, Steffen Rechner†, and Matthias Müller-Hannemann

Institute of Computer Science
Martin Luther University Halle-Wittenberg, Germany

{erbert,rechner,muellerh}@informatik.uni-halle.de

October 21, 2018

Abstract

A basic task in bioinformatics is the counting of k-mers in genome strings. The *k-mer counting problem* is to build a histogram of all substrings of length k in a given genome sequence. We present the open source k-mer counting software *Gerbil* that has been designed for the efficient counting of k-mers for $k \geq 32$. Given the technology trend towards long reads of next-generation sequencers, support for large k becomes increasingly important. While existing k-mer counting tools suffer from excessive memory resource consumption or degrading performance for large k, *Gerbil* is able to efficiently support large k without much loss of performance. Our software implements a two-disk approach. In the first step, DNA reads are loaded from disk and distributed to temporary files that are stored at a working disk. In a second step, the temporary files are read again, split into k-mers and counted via a hash table approach. In addition, *Gerbil* can optionally use GPUs to accelerate the counting step. For large k, we outperform state-of-the-art open source k-mer counting tools for large genome data sets.

1 Introduction

The counting of k-mers in large amounts of reads is a common task in bioinformatics. The problem is to count the occurrences of all k-long substrings in a large amount of sequencing reads. Its most prominent application is de novo assembly of genome sequences. Although building a histogram of k-mers seems to be quite a simple task from an algorithmic point of view, it has attracted a considerably amount of attention in recent years. In fact, the counting of k-mers becomes a challenging problem for large instances, if it is to be both resource- and time-efficient and therefore makes it an interesting object of study for algorithm engineering. Existing tools for k-mer counting are often optimized for $k < 32$ and lack good performance for larger k. Recent advances in technology towards larger read lengths are leading to the quest to cope with values of k exceeding 32. Studies elaborating on the optimal choice for the value of k recommend for various applications relatively high values [1,13]. In particular, working with long sequencing reads helps to improve accuracy and contig assembly (with k values in the hundreds) [11]. In this paper, we develop a tool with a high performance for such large values of k.

*A short version of this paper will appear in the proceedings of WABI 2016.

†Corresponding author.

1.1 Related Work

Among the first software tools that succeeded in counting the k-mers of large genome data sets was Jellyfish [5], which uses a lock-free hash table that allows parallel insertion. In the following years, several tools were published, successively reducing running time and required memory. BFCounter [6] uses Bloom filters for k-mer counting to filter out rarely occurring k-mers stemming from sequencing errors. Other tools like DSK [7] and KMC [2] exploit a two-disk architecture and aim at reducing expensive IO operations. Turtle [10] replaces a standard Bloom filter by a cache-efficient counterpart. MSPKmerCounter [4] introduces the concept of minimizers to the k-mer counting, thus further optimizing the disk-based approach. The minimizer approach was later on refined to signatures within KMC2 [3]. Up to now, the two most efficient open source software tools have been KMC2 and DSK. KMC2 uses a sorting based counting approach that has been optimized for $k < 32$. However, its performance drops when k grows larger. Instead, DSK uses a single large hash table and is therefore efficient for large k (but does not support $k > 127$). However, for small k, it is clearly slower than KMC2. To the best of our knowledge, the only existing approach that uses GPUs for counting k-mers is the work by Suzuki et al. [12].

1.2 Contribution

In this article we present the open source k-mer counting tool *Gerbil*. Our software is the result of an extensive process of algorithm engineering that tried to bring together the best ideas from the literature. The result is a k-mer counting tool that is both time efficient and memory frugal.[1] In addition, *Gerbil* can optionally use GPUs to accelerate the counting step. It outperforms its strongest competitors both in efficiency and resource consumption significantly. For large values of k, it reduces the runtime by up to a factor of four. The software is written in C++ and CUDA and is freely available at `https://github.com/uni-halle/gerbil` under MIT license.

In the next section we describe the general algorithmic work flow of *Gerbil*. Thereafter, in Section 3, we focus on algorithm engineering aspects that proved essential for high performance and describe details, like the integration of a GPU into the counting process. In Section 4, we evaluate *Gerbil*'s performance in a set of experiments and compare it with those of KMC2 and DSK. We conclude this article by a short summary and a glance on future work.

2 Methods

Gerbil is divided into two phases: (1) Distribution and (2) Counting. In this section, we give a high-level description of *Gerbil*'s work flow.

2.1 Distribution

Whole genome data sets typically do not fit into the main memory. Hence, it is necessary to split the input data into a couple of smaller temporary files. *Gerbil* uses a two-disk approach that is similar to those of most contemporary k-mer counting tools [3, 4, 7]. The first disk contains the input read data and is used to store the counted k-mer values. We call this disk input/output-disk. The second disk, which we call working disk, is used to store temporary files. The key idea is to assure that the temporary files partition the input reads in such a way, that all occurrences of a certain k-mer are stored in the same temporary file. This way, one can simply count the k-mers of the temporary files independently of each other, with small main memory requirements. To split the genome data into temporary files, we make use of the *minimizer* approach that has been proposed by [9] and later on refined by [3]. A genome sequence can be decomposed into a number of overlapping *super-mers*. Each super-mer is a substring of maximal length such that all k-mers on that substring share the same minimizer. Hereby, a minimizer of a k-mer is defined as its lexicographically

[1] The tool is named *Gerbil* because of its modest resource requirements, which it has in common with the name-giving mammal.

```
CAAGAACAGTG
CAAG                    1. CAAGA
 AAGA
  AGAA                  2. AGAA
   GAAC                 3. GAACA
    AACA
     ACAG               4. ACAG
      CAGT              5. CAGTG
       AGTG
```

Figure 1: Minimizers and super-mers of the DNA string `CAAGAACAGTG`. Here, $k = 4$ and $m = 3$. For each k-mer, the bold part is its minimizer. The example uses the lexicographic ordering on 3-mers based on $A < C < G < T$. The sequence is divided into the five super-mers `CAAGA`, `AGAA`, `GAACA`, `ACAG`, and `CAGTG` that would be stored in temporary files.

smallest substring of a fixed length $m < k$ with respect to some total ordering on strings of length m. See Fig. 1 for an example. It suffices to partition the set of super-mers into different temporary files to achieve a partitioning of all different k-mers [3].

2.2 Counting

The counting of k-mers is typically done by one of two approaches: Sorting and Compressing [3] or using a hash table with k-mers as keys and counters as values [5, 7]. The efficiency of the sorting approach typically relies on the sorting algorithm Radix Sort, whose running time increases with the length of k-mers. Since we aim at high efficiency for large k, we decided to implement the hash table approach. Therefore, we use a specialized hash table with k-mers as keys and counters as values. We use a hash table that implements open addressing and solves collisions via double hashing. Alg. 1 shows a high level description of the insertion method.

Algorithm 1 Insert the k-mer x into the hash table t.

Input: Hash Table t, k-mer x, Maximal number of trials θ.

```
i ← 0
while i < θ do                          /* while max number of trials not reached */
    p ← hash(x, i)
    if t_p = (x, c) then                /* matching k-mer detected */
        t_p ← (x, c + 1)
        return
    else if t_p is empty then           /* empty table entry */
        t_p ← (x, 1)
        return
    else                                /* entry locked by another k-mer */
        i ← i + 1
    end if
end while
Start emergency mechanism.
```

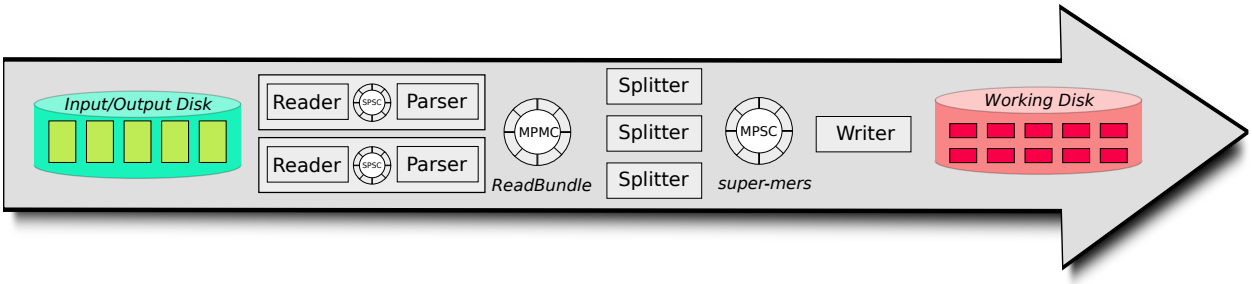

Figure 2: Work flow of Phase One.

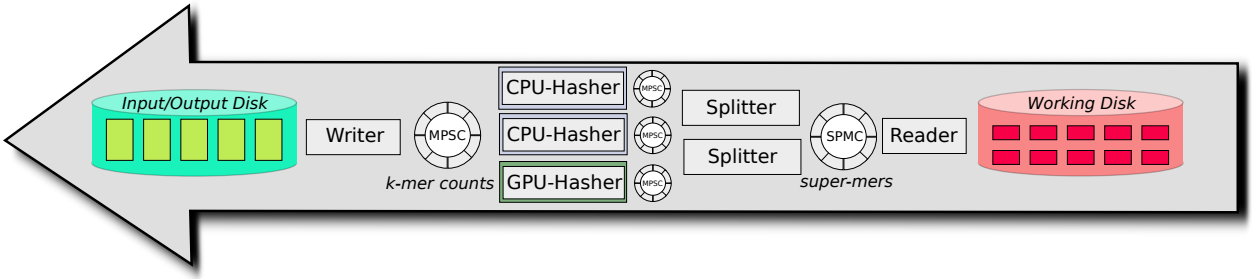

Figure 3: Work flow of Phase Two.

2.3 Work Flow

Although the following description of the main process is sequential, all of the steps are interleaved and therefore executed in parallel. This is done by a classical pipeline architecture. Each output of a step makes the input of the next. We use ring buffers to connect the steps of the pipeline. Such buffers are specialized for all combinations of single (S)/multiple (M) producers (P) and single (S)/multiple (M) consumers (C). The actual number of parallel threads depend on the system and is determined by the software at runtime to achieve optimal memory throughput.

2.3.1 Phase One: Distribution

The goal of the first phase is to split the input data into a number of temporary files. Fig. 2 visualizes the first phase.

1. A group of reader threads read the genome reads from the input disk into the main memory. For compressed input, these threads also decompress it.
2. A second group of parser threads convert the read data from the input format into an internal read bundle format.
3. A group of splitter threads compute the minimizers of the reads. All subsequent substrings of a read that share the same minimizer are stored as a super-mer into an output buffer.
4. A single writer thread stores the output buffers to a variable number of temporary files at the working disk.

2.3.2 Phase Two: Counting

After the first phase has been completed, the temporary files are sequentially re-read from working disk and processed in the following manner (see Fig. 3).

1. A single reader thread reads the super-mers of a temporary file and stores them in main memory.
2. A group of threads split the super-mers into k-mers. Each k-mer is distributed to one of multiple hasher threads by using a hash function on each k-mer. This ensures that multiple occurrences of the same k-mer are assigned to the same hasher thread and allows the distribution of separated hash tables to different memory spaces.
3. A group of hasher threads insert the k-mers into their thread-own hash tables. After a temporary file has been completely processed, each hasher thread sends the content of its hash table to an output buffer.
4. A single writer thread writes from the output buffer to the output disk.

2.4 DNA Sequence Handling

2.4.1 Undetermined bases

DNA reads typically contain bases that could not been identified correctly during the sequencing process. Usually, such bases are marked N in FASTQ input files. In accordance with established k-mer counting tools, we ignore all k-mers that contain an undetermined base.

2.4.2 Reverse-Complement

Since DNA is organized in double helix form, each k-mer $x \in \{A, C, G, T\}^k$ corresponds to its *reverse-complement* that is defined by reversing x and replacing $A \Leftrightarrow T$ and $C \Leftrightarrow G$. Thus, the k-mer $ACCG$ corresponds to $CGGT$. Many applications do not distinguish between a k-mer and its reverse-complement. Thus, each occurrence of $ACCG$ and $CGGT$ is counted as occurrences of their unique *canonical* representation. *Gerbil* uses the lexicographically smaller k-mer as canonical representation. The use of reverse complement normalization can be turned off by command flag.

3 Implementation Details

We now want to point out several details on the algorithm engineering process that were essential to gain high performance.

3.1 Total ordering on minimizers

The choice of a total ordering has large effects on the size of temporary files and thus, also on the performance. To find a good total ordering, we have to balance various aspects. On the one hand, the total number of resulting super-mers are to be minimized to reduce the total size of disk memory that is needed by temporary files. On the other hand, the maximal number of distinct k-mers that share the same minimizer should not be too large since we want an approximately uniform distribution of k-mers to the temporary files. An "ideal" total ordering would have both a large total number of super-mers and a small maximal number of distinct k-mers per minimizer. Since these requirements contradict each other, we experimentally evaluated the pros and cons of various ordering strategies.

CGAT The lexicographic ordering of minimizers based on $C < G < A < T$.

Roberts et al. [8] They propose the lexicographic ordering of minimizers with respect to $C < A < T < G$. Furthermore, within the minimizer computation all bases at even positions are to be replaced by their reverse complement. Thus, rare minimizers like $CGCGCG$ are preferred.

KMC2 The ordering that is proposed by [3] is a lexicographic ordering with $A < C < G < T$ and some built-in exceptions to eliminate the large number of minimizers that start with AAA or ACA.

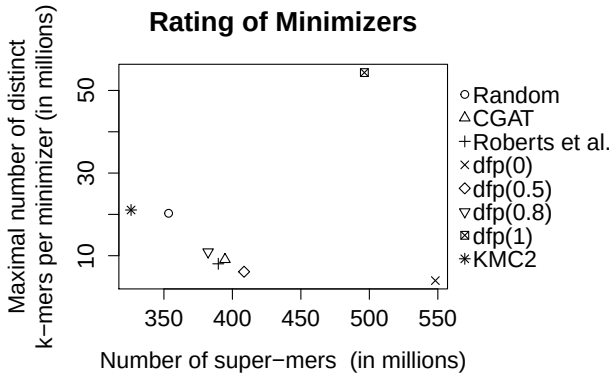

Figure 4: Evaluation of various total ordering strategies for minimizers (*F Vesca*, $m = 6$, $k = 28$). Strategy dfp(p) has been tested with $p \in \{0, 0.5, 0.8, 1\}$.

Random A random order of all string of fixed length m is unlikely to have both a small number of super-mers and a highly imbalanced distribution of distinct k-mers. It is simple to establish, since we do not need frequency samples or further assumptions about the distribution of minimizers.

Distance from Pivot (dfp(p)) To explain this strategy, consider the following observations: Ascendingly sorting the minimizers by their frequency favors rare minimizers. As a consequence, the maximal number of distinct k-mers per minimizer is small. However, the total number of super-mers can be very large. Similarly, an descendingly sorted ordering results in quite the opposite effect. To find a compromise between both extremes, we initially sort the set of minimizers by their frequency. Since the frequencies depend on the data set, we approximate them by taking samples during runtime. We fix a pivot factor $0 \leq p \leq 1$ and re-sort the minimizers by the absolute difference of their initial position to the pivot position $4^m p$. The result is an ordering that does neither prefer very rare nor very common minimizers and therefore makes a good compromise.

3.1.1 Evaluation

See Fig. 4 for a rating of each strategy. The value on the x-axis corresponds to the expected temporary disk memory, whereas the value on the y-axis is correlated with the maximal main memory consumption of our program. A perfect strategy would be located at the bottom left corner. Several strategies seem to be reasonable choices. We evaluated each strategy and found that a small number of super-mers is more important than a small maximal number of k-mer per minimizer for most data sets. As a result, we confirm that the total ordering that is already been used by KMC2 is a good choice for most data sets. Therefore, *Gerbil* uses the strategy from KMC2 for its ranking of minimizers.

3.2 Length of miminizers

The length m of miminizers is a parameter that has to be chosen with care. However, we can consider a basic rule: The larger m is chosen, the less likely it becomes that consecutive k-mers share the same minimizer. Therefore, the number of super-mers decreases with growing m. An advantage of a smaller number of super-mers is that the set of super-mers can be distributed to the temporary files more uniformly, which results in temporary files of approximately uniform size. However, a major drawback of a large number of super-mers is the increased total size of temporary files. Thus, a small m results in a better data compression.

In our experiments, we found that choosing minimizer length $m = 7$ is most efficient for the data sets.

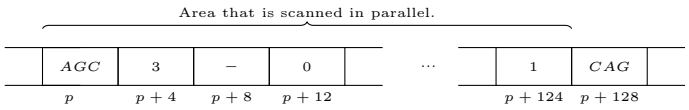

Figure 5: GPU memory access pattern. The figure shows the memory area that is being scanned while probing a hash table entry that is stored at memory address p. In this example, $k = 3$ and each table entry needs four bytes for the key and four bytes for the counter. Therefore, 16 entries can be loaded from global memory within one step and are scanned in parallel.

3.3 GPU Integration

To integrate one or more GPUs into the process of k-mer counting, several problems have to be dealt with. Typically, a GPU performs well only if it deals with data in a parallel manner. In addition, memory bound tasks (i. e. tasks that do not require a lot of arithmetic operations) like the counting of k-mers require a carefully chosen memory access pattern to minimize the number of the accesses to the GPU's global memory. We decided to transfer the hash table based counting approach to the GPU.

3.3.1 GPU Hash Tables

When compiled and executed with GPU support, *Gerbil* automatically detects CUDA capable GPUs. For each GPU, *Gerbil* replaces a CPU hasher thread by a GPU hasher thread which maintains its own hash table in GPU memory. Each GPU hash table is similar in function to a traditional hash table. However, unlike the traditional approach, we add a large number of k-mers in parallel. Therefore, the insertion procedure is slightly changed.

First, a bundle of several thousand k-mers is copied to the GPU global memory space. Afterwards, we launch a large number of CUDA blocks, each consisting of 32 threads. Each block sequentially inserts a few k-mers into the GPU hash table. Since with increasing running time, it becomes more and more probable to find a mismatch when probing a hash table position, we additionally scan adjacent table positions in a range of 128 bytes when probing a hash table entry (see Fig. 5). Due to the architecture of a GPU, this can be done within the same global memory access. Thus, we scan up to 16 table entries in parallel, thereby reducing the number of accesses to a GPU's global memory. In addition, the total number of probing operations is drastically reduced. In particular, by probing just one entry after the other, 90.37% of all 28-mers of the *F Vesca* data set (Table 1) can be inserted at first probing and no 28-mer needed more than 29 probings. In contrast, through scanning of adjacent table entries, 99.94% of all 28-mers could be inserted at the first trial and no 28-mer needed more than seven probing operations. To eliminate race conditions between CUDA blocks, we synchronize the probing of the hash table by using atomic operations to lock and unlock hash table entries. Since such operations are efficiently implemented in hardware, a large number of CUDA blocks can be executed in parallel.

3.3.2 Load Balancing

We dynamically balance the amount of k-mers that are assigned to the various CPU and GPU hasher threads. Therefore, we constantly measure the throughput of each hasher thread, i. e. the CPU-time needed to insert a certain number of k-mers. Whenever a new temporary file is loaded from disk, we rebalance the number of k-mers that are assigned to each hasher thread, considering the throughput and capacity of each hash table. By that, we automatically determine a good division of labour between CPU and GPU hasher threads without the need of careful hand-tuning.

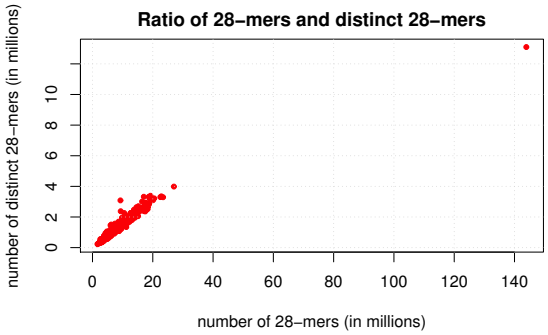

(a) Number of 28-mers and number of distinct 28-mers in the 512 temporary files that have been created while processing the *F Vesca* data set. Each point corresponds to a temporary file. Here, the KMC2 minimizer ordering did not succeed in creating uniformly sized temporary files since a single file contains far more 28-mers than the other 511 files.

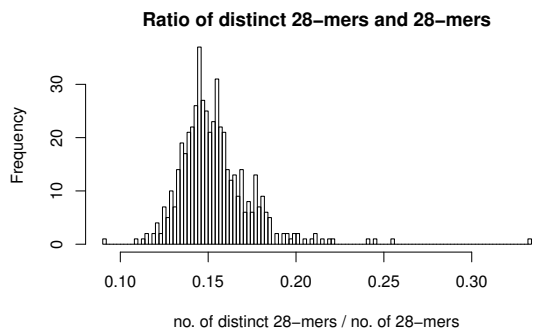

(b) Dividing the number of 28-mers in each file by number of its distinct 28-mers leads to a ratio that is used to determine the size of the hash tables. A ratio between 0.15 and 0.2 is a proper choice for the *F Vesca* data set.

Figure 6: Estimation of hash table sizes.

3.4 Hash Table Details

3.4.1 Estimating table sizes

We aim at estimating the expected size of each hash table as closely as possible to save main memory. We do so since reduced memory consumption leaves more memory to the operating system that can be used as cache when writing temporary files. Therefore, we approximate the number of expected distinct k-mers in each temporary file. We use a simple approximation mechanism that predicts the number of distinct k-mers in a file by multiplying the number of k-mers in each file with a constant that has been determined experimentally (see Fig. 6). Since this ratio depends on properties of the data set, we dynamically adjust the ratio during runtime.

3.4.2 Hash Function

To insert a k-mer into the hash table, we use a hash function that implements double hashing.

```
/**
 * Hash function based on byte-wise interpretation of k-mers.
 * @param kmer: a kmer that shall be inserted
 * @param trial: the number of probings
 */
uint32_t hash(const Kmer<k>& kmer, const uint32_t trial) {
    uint8_t* keyBytes = (uint8_t*) &kmer;   // interpret kmer as byte array
    uint32_t h1 = 0; // value of first hash function
    uint32_t h2 = 0; // value of second hash function
    for (uint32_t i = 0; i < sizeof(kmer); i++) {
        h1 = 31 * h1 + keyBytes[i];
        h2 = 37 * h2 + ~keyBytes[i];
    }
    return h1 + trial * h2;
}
```

Table 1: Data Sets.

Data Set	Format	Size (GB)	Read Length	28-mers	distinct 28-mers
F Vesca	FASTQ	10.2	353	4 134 078 256	632 436 468
M Balbisiana	FASTQ	98.6	100	20 531 572 597	965 691 662
G Gallus	FASTQ	115.9	101	25 337 974 831	2 727 529 829
H Sapiens	FASTQ	223.3	100	62 739 461 708	6 336 805 684
H Sapiens 2	FASTQ	339.5	101	98 892 620 173	6 634 382 141
GRCh38	FASTA	100.0	1000	97 300 000 000	1 802 953 276

Table 2: Test Systems.

	System One	System Two
CPU	Intel Core-i5 2550k (4 cores)	Intel Xeon(R) E3-1231 v3 (8 cores)
RAM	16 GB DDR3	32 GB DDR3
GPU	GeForce GTX 970	GeForce GTX TITAN X GeForce GTX 970
Working-Disk	256 GB Crucial M550	2x Samsung 850 EVO 500 GB (RAID-0)
Free disk space	128 GB	1000 GB
OS	Ubuntu 14.04 LTS	
In/Out-Disk	Transcend StoreJet 35T3 USB 3.0 (External HDD)	

3.4.3 Probing Strategy

As a general strategy we use double hashing. We stop the probing of the hash table after a constant number of trials. Therefore, it is possible that k-mers could not be inserted into a hash table. For that reason, *Gerbil* has a built-in emergency mechanism that handles such k-mers to prevent them from getting lost. Hereby, CPU and GPU hasher threads have different strategies. CPU hasher threads store such k-mers in an additional temporary file, which is processed after the work with the current temporary file has been completed. In contrast, GPU hasher threads use part of free GPU memory to sequentially store those k-mers that could not be inserted. After all k-mers of a temporary file have been processed, the k-mers in this area are counted via a sorting and compression approach. However, it is still possible to exceed the available GPU memory. In such a case, we copy the whole amount of k-mers in that area back to main memory and store them in a temporary file, similar to the CPU emergency handling. Such an operation is very costly. However, we have never observed a single GPU error handling and only few executions of CPU error handling when processing real world data sets.

4 Results

We tested our implementation in a set of experiments, using the same instances as Deorowicz et al. [3] (see Table 1). For each data set we counted all k-mers for $k = 28, 40, 56$, and 65 and compared *Gerbil*'s running time with those of KMC2 in version 2.3.0 and DSK in version 2.0.7. In addition, we used a synthesized test set *GRCh38*, created from Genome Reference Consortium Human Reference 38 (GCA_000001405.2), from which we uniformly sampled k-mers of size 1000. The purpose of this data set is to have longer reads allowing to test the performance for larger values of k. To judge performance on various types of hardware, we executed the experiments on two different desktop computers. See Table 2 for details about the hardware configuration of the test systems. Table 3 and Fig. 7 show the results of the performance evaluation. We want to point out several interesting observations.

- *Gerbil* with GPU support (*gGerbil*) is the most efficient tool in almost all cases. Exceptions occur for small $k = 28$, where the sorting based approach KMC2 is sometimes slightly more efficient.

Table 3: Running times in the format mm:ss (the best performing in bold). Each entry is the average over three runs. Missing running times for DSK are due to insufficient disk space. The label 'gGerbil' stands for *Gerbil* with activated GPU mode. Instead, standard 'Gerbil' does not use any GPU.

		System One				**System Two**			
Data Set	k	Gerbil	gGerbil	KMC2	DSK	Gerbil	gGerbil	KMC2	DSK
F Vesca	28	02:08	**01:40**	02:01	03:00	01:36	**01:18**	01:32	02:05
	40	02:34	**01:53**	03:03	04:14	02:01	**01:38**	02:12	02:52
	56	02:58	**01:53**	03:19	03:55	02:25	**01:39**	02:30	02:50
	65	03:05	**01:59**	04:34	05:23	02:16	**01:42**	03:35	03:37
M Balbisiana	28	13:37	**11:42**	12:54	14:49	11:17	**10:07**	10:50	11:06
	40	13:48	**12:24**	16:15	16:12	11:46	**10:59**	13:46	12:26
	56	12:46	**11:36**	16:06	14:56	10:50	**10:18**	13:36	11:44
	65	12:32	**11:28**	18:33	15:52	10:46	**10:16**	15:47	12:34
G Gallus	28	18:41	**14:25**	15:39	26:54	15:47	**12:31**	13:10	21:00
	40	19:55	**16:00**	19:44	29:42	16:29	**14:10**	16:49	23:48
	56	18:12	**14:48**	19:48	24:11	15:38	**13:12**	16:48	19:59
	65	18:27	**15:22**	22:49	26:50	15:41	**13:08**	19:25	21:33
H Sapiens	28	41:10	**30:04**	32:18	-	33:26	**25:16**	26:44	50:15
	40	45:02	**35:52**	43:19	-	35:20	**29:00**	35:59	54:21
	56	39:47	**33:21**	42:53	-	32:21	**26:46**	35:25	45:32
	65	38:09	**35:32**	51:23	-	32:09	**26:27**	42:19	47:50
H Sapiens 2	28	65:33	49:41	**49:17**	-	53:40	**39:24**	41:47	76:50
	40	72:06	**66:04**	70:33	-	57:03	**46:00**	57:02	83:59
	56	64:00	**60:27**	69:58	-	51:34	**42:15**	56:28	72:35
	65	**61:05**	64:44	87:24	-	51:16	**41:30**	68:10	78:13

- Fig. 7a shows that Gerbil starts outperforming KMC2 at $k \approx 36$. Interestingly, the running time of each tool decreases with growing k. This can be explained by the small read length of the *G Gallus* data set. In addition, one can observe the erratic increase of running time near $k = 32$ and $k = 64$ for all tools, due to a change of the internal k-mer representation.
- When k grows, KMC2 becomes more and more inefficient, while *Gerbil* stays efficient. When counting the 200-mers in the *GRCh38* data set, KMC2 did not finish within 20 hours, whereas *Gerbil* required only 98 minutes (Fig. 7b). The running time of DSK grows similarly fast as that of KMC2. Recall that DSK does not support values of $k > 127$.
- For small k, the use of a GPU decreases the running time by a significant amount of time. However, with growing k, the data structure that stores k-mers grows larger. Therefore, the number of table entries that can be scanned in parallel decreases. Experimentally, we found that the GPU induced speedup vanishes when k exceeds 150.

We gain some additional interesting insights when we take a closer look into Table 4 that shows detailed information on running time and memory usage.

- The use of a GPU accelerates *Gerbil*'s second phase by up to a factor of two, whereas the additional speedup given by a second GPU is only moderate.
- All tools were called with an option that sets the maximal memory size to 14 GB on Test System One and 30 GB on Test System Two. However, *Gerbil* typically uses much less memory due to its dynamic prediction of the hash table size. In contrast, both KMC2 and DSK use more main memory.
- *Gerbil*'s disk usage is comparable to KMC2's disk usage, whereas the disk usage of DSK is much larger.
- *Gerbil*'s frugal use of disk- *and* main memory is a main reason for its high performance. The use of little main memory gives the operating system opportunity to use the remaining main memory for

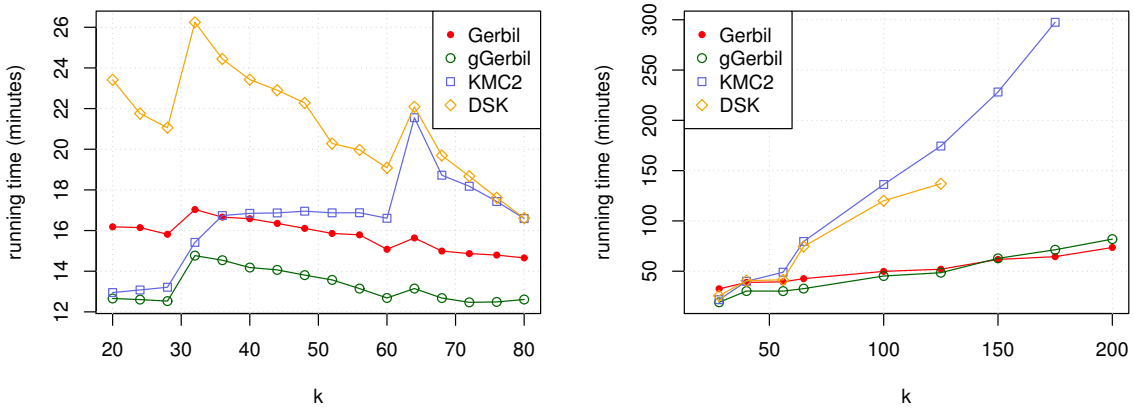

(a) Running times for $20 \leq k \leq 80$ of *G Gallus* data set.

(b) Running times for $28 \leq k \leq 200$ of *GRCh38* data set.

Figure 7: Running time for various k. (Test System Two)

Table 4: Detailed running times (in format mm:ss) and maximal main memory and disk space consumption (in GB) for the *G Gallus* instance. Each entry is the average of three runs.

		System One				**System Two**			
	k	Gerbil	gGerbil	KMC	DSK	Gerbil	gGerbil	KMC	DSK
Phase 1	28	10:08	10:06	10:51	10:22	09:46	09:43	09:52	09:30
Phase 2	28	08:32	04:19	04:46	16:00	06:01	02:47	03:16	11:01
Main Memory	28	2.36	1.77	14.28	15.28	2.20	2.01	26.99	16.69
Disk Space	28	23.66	23.66	24.86	37.30	23.66	23.66	24.86	37.30
Phase 1	56	10:07	10:06	10:40	10:26	09:47	09:43	09:47	09:30
Phase 2	56	08:05	04:42	09:08	13:13	05:50	03:28	06:59	10:00
Main Memory	56	4.24	3.20	14.29	15.00	4.00	3.40	26.98	14.78
Disk Space	56	16.25	16.25	17.02	57.20	16.25	16.25	17.02	57.20

buffering disk operations. A small disk space consumption is essential since disk operations are far more expensive than the actual counting.

- We can compare the actual running time with the theoretically minimal running time that is given by hardware constraints (hereby done with Test System Two). We find that the running time of phase one is bounded by the read rate of the input disk. Thus, a further speedup of the first phase can only be achieved by reading compressed input files. In contrast, there is potential speedup in phase two: In GPU mode, Gerbil's throughput in phase two is just about 61% of the possible throughput given by write rate of the output disk. In non-GPU mode it is about 35%. Interestingly, the throughput of the working disk is not a critical component at our test system.

5 Conclusion

We introduced the k-mer counting software *Gerbil* that uses a hash table based approach for the counting of k-mers. For large k, a use case that becomes important for long reads, we are able to clearly outperform the state-of-the-art open source k-mer counting tools, while using significantly less resources. We showed that *Gerbil*'s running time can be accelerated by the use of GPUs. However, since this only affects the second phase, the overall additional speedup is only very moderate. As future work, we plan to evaluate

strategies to use GPUs to accelerate also the first phase. Another option for further speed-up would be to give up exactness by using Bloom filters.

References

[1] Chikhi, R., Medvedev, P.: Informed and automated k-mer size selection for genome assembly. Bioinformatics 30(1), 31–37 (2014)

[2] Deorowicz, S., Debudaj-Grabysz, A., Grabowski, S.: Disk-based k-mer counting on a PC. BMC Bioinformatics 14(1), 1–12 (2013)

[3] Deorowicz, S., Kokot, M., Grabowski, S., Debudaj-Grabysz, A.: KMC 2: fast and resource-frugal k-mer counting. Bioinformatics 31(10), 1569–1576 (2015)

[4] Li, Y., et al.: MSPKmerCounter: a fast and memory efficient approach for k-mer counting. arXiv preprint arXiv:1505.06550 (2015)

[5] Marçais, G., Kingsford, C.: A fast, lock-free approach for efficient parallel counting of occurrences of k-mers. Bioinformatics 27(6), 764–770 (2011)

[6] Melsted, P., Pritchard, J.K.: Efficient counting of k-mers in DNA sequences using a bloom filter. BMC Bioinformatics 12(1), 1–7 (2011)

[7] Rizk, G., Lavenier, D., Chikhi, R.: DSK: k-mer counting with very low memory usage. Bioinformatics 29(5), 652–3 (2013)

[8] Roberts, M., Hayes, W., Hunt, B.R., Mount, S.M., Yorke, J.A.: Reducing storage requirements for biological sequence comparison. Bioinformatics 20(18), 3363–3369 (2004)

[9] Roberts, M., Hunt, B.R., Yorke, J.A., Bolanos, R.A., Delcher, A.L.: A preprocessor for shotgun assembly of large genomes. Journal of Computational Biology 11(4), 734–752 (2004)

[10] Roy, R.S., Bhattacharya, D., Schliep, A.: Turtle: Identifying frequent k-mers with cache-efficient algorithms. Bioinformatics 30(14), 1950–7 (2014)

[11] Sameith, K., Roscito, J.G., Hiller, M.: Iterative error correction of long sequencing reads maximizes accuracy and improves contig assembly. Briefings in Bioinformatics (2016), `http://dx.doi.org/10.1093/bib/bbw003`

[12] Shuji Suzuki, Masanori Kakuta, T.I.Y.A.: Accelerating identification of frequent k-mers in DNA sequences with GPU. In: GTC 2014 (2014)

[13] Xavier, B.B., Sabirova, J., Pieter, M., Hernalsteens, J.P., de Greve, H., Goossens, H., Malhotra-Kumar, S.: Employing whole genome mapping for optimal de novo assembly of bacterial genomes. BMC Research Notes 7(1), 1–4 (2014)

Appendix

A Command Line Parameters

Gerbil can be controlled by several command line arguments and flags.

Command	Description	Default
`-k length`	Set the value of k, i. e. the length of k-mers to be counted. Supported k range from 8 to 479.	28
`-m length`	Set the length m of minimizers.	auto
`-e size`	Restrict the maximal size of main memory in `MB` or `GB` that *Gerbil* is allowed to use.	auto
`-f number`	Set the number of temporary files.	512
`-t number`	Set the maximal number of parallel threads to use.	auto
`-l count`	Set the minimal occurrence of a k-mer to be outputted.	3
`-i`	Enable additional output.	
`-g`	Enable GPU mode. *Gerbil* will automatically detect CUDA-capable devices and will use them for counting in the second phase.	
`-v`	Show version number.	
`-d`	Disable normalization of k-mers. If normalization is disabled, a k-mer and its reverse complement are considered as different k-mers. If normalization is enabled, we map both k-mer and its reverse complement to the same k-mer.	
`-s`	Perform a system check and display information about your system.	
`-x 1`	Stop execution after Phase One. Do not remove temporary files and a `binStatFile` (with statistical information). Watch out: no `output` allowed.	
`-x 2`	Only execute Phase Two. Requires temporary files and the `binStatFile`.	
`-x b`	Do not remove the `binStatFile`.	
`-x h`	Create a histogram of k-mers in a human readable format in output directory.	

B Input Formats

Gerbil supports the following input formats of genome read data: FASTQ, FASTA, staden, as well as compressed files of these formats. To process multiple files, it also supports textfiles with paths to one or more input files.

C Output Format

Gerbil uses an output format that is easy to parse and requires little space. The counter of each occuring k-mer is stored in binary form, followed by the corresponding byte-encoded k-mer. Each four bases of a k-mer are encoded in one single byte. We encode A with 00, C with 01, G with 10 and T with 11. Most counters of k-meres are slightly smaller than the coverage of the genome data. We exploit this property by using only one byte for counters less than 255. A counter greater than or equal to 255 is encoded in five bytes.

In the latter case, all bits of the first byte are set to 1. The remaining four bytes contain the counter in a conventional 32-bit unsigned integer. Examples (X is undefined):

- `67 AACGTG` ⇒ `01000011 00000110 1110XXXX`
- `345 TGGATC` ⇒ `11111111 00000000 00000000 00000001 01011001 11101000 1101XXXX`

When called with command line argument `-x h`, *Gerbil* additionally creates a human readable csv file including the same, but uncompressed information. This option should only be used for very small data sets.

D Datasets

The data sets were downloaded from the following URL's.

F Vesca

ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA020/SRA020125/SRX030575/SRR072005.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA020/SRA020125/SRX030576/SRR072006.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA020/SRA020125/SRX030576/SRR072007.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA020/SRA020125/SRX030577/SRR072008.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA020/SRA020125/SRX030577/SRR072009.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA020/SRA020125/SRX030575/SRR072010.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA020/SRA020125/SRX030575/SRR072011.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA020/SRA020125/SRX030575/SRR072012.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA020/SRA020125/SRX030578/SRR072013.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA020/SRA020125/SRX030578/SRR072014.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA020/SRA020125/SRX030578/SRR072029.fastq.bz2

G Gallus

ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA030/SRA030308/SRX043656/SRR105788_1.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA030/SRA030308/SRX043656/SRR105788_2.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA030/SRA030309/SRX043656/SRR105789_1.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA030/SRA030309/SRX043656/SRR105789_2.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA030/SRA030312/SRX043656/SRR105792_1.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA030/SRA030312/SRX043656/SRR105792_2.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA030/SRA030314/SRX043656/SRR105794.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA030/SRA030314/SRX043656/SRR105794_1.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA030/SRA030314/SRX043656/SRR105794_2.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA036/SRA036382/SRX043656/SRR197985.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA036/SRA036382/SRX043656/SRR197985_1.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA036/SRA036382/SRX043656/SRR197985_2.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA036/SRA036383/SRX043656/SRR197986.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA036/SRA036383/SRX043656/SRR197986_1.fastq.bz2
ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA036/SRA036383/SRX043656/SRR197986_2.fastq.bz2

M Balbisiana

ftp://ftp.ddbj.nig.ac.jp/ddbj_database/dra/fastq/SRA098/SRA098922/SRX339427/SRR956987.fastq.bz2

H Sapiens

ftp://ftp.1000genomes.ebi.ac.uk/vol1/ftp/phase3/data/HG02057/sequence_read/SRR359301.filt.fastq.gz
ftp://ftp.1000genomes.ebi.ac.uk/vol1/ftp/phase3/data/HG02057/sequence_read/SRR359301_1.filt.fastq.gz
ftp://ftp.1000genomes.ebi.ac.uk/vol1/ftp/phase3/data/HG02057/sequence_read/SRR359301_2.filt.fastq.gz
ftp://ftp.1000genomes.ebi.ac.uk/vol1/ftp/phase3/data/HG02057/sequence_read/SRR360755.filt.fastq.gz
ftp://ftp.1000genomes.ebi.ac.uk/vol1/ftp/phase3/data/HG02057/sequence_read/SRR360755_1.filt.fastq.gz
ftp://ftp.1000genomes.ebi.ac.uk/vol1/ftp/phase3/data/HG02057/sequence_read/SRR360755_2.filt.fastq.gz

H Sapiens 2

http://www.ebi.ac.uk/ena/data/view/ERA015743/

GRCh38

http://hgdownload.cse.ucsc.edu/goldenPath/hg38/bigZips/hg38.fa.gz