

SlopMap: A Software Application Tool for Quick and Flexible Identification of Similar Sequences Using Exact K-Mer Matching

Ilya Y Zhbannikov^{1*}, Samuel S Hunter^{1,2}, Matthew L Settles^{1,2} and James A Foster^{1,2,3}

¹Institute for Bioinformatics and Evolutionary Studies, University of Idaho, USA

²The Institute of Bioinformatics and Evolutionary Studies (IBEST), University of Idaho, USA

³Department of Biological Sciences, University of Idaho, USA

Abstract

With the advent of Next-Generation (NG) sequencing, it has become possible to sequence an entire genome quickly and inexpensively. However, in some experiments one only needs to extract and assembly a portion of the sequence reads, for example when performing transcriptome studies, sequencing mitochondrial genomes, or characterizing exomes. With the raw DNA-library of a complete genome it would appear to be a trivial problem to identify reads of interest. But it is not always easy to incorporate well-known tools such as BLAST, BLAT, Bowtie, and SOAP directly into a bioinformatics pipelines before the assembly stage, either due to in-compatibility with the assembler's file inputs, or because it is desirable to incorporate information that must be extracted separately. For example, in order to incorporate flowgrams from a Roche 454 sequencer into the Newbler assembler it is necessary to first extract them from the original SFF files.

We present SlopMap, a bioinformatics software utility which allows rapid identification similar to provided target sequences from either Roche 454 or Illumina DNA library. With a simple and intuitive command-line interface along with file output formats compatible with assembly programs, SlopMap can be directly embedded in biological data processing pipeline without any additional programming work. In addition, SlopMap preserves flowgram information needed for Roche 454 assembler.

Introduction

New methodologies enabled by Next Generation Sequencing (NGS) that are of particular interest to us include transcriptome analysis for RNA research [1] and mitochondrial sequencing from exome data [2]. Such applications include those in which the researcher is interested in assembling only specific content within a genome of interest, using a set of targets to initialize the assembly process. It may seem trivial to identify the reads of interest among those produced by NGS hardware, using well-known general-purpose alignment or mapping tools such as Blat [3], Bowtie 2 [4], BWA [5], and SOAP [6]. But even an efficient tool may be difficult to incorporate directly into a bioinformatics pipeline before the assembly stage, since it may be necessary to convert data to a different file format. For example, existing mappers usually use the SAM/BAM [7] file format as output. None use SFF format files [8] as both input and output, and none but Bowtie2 support FastQ output, and it is supported only in a limited sense.

Moreover, it is difficult to use existing mapping software tools when it necessary to establish a similarity threshold, i.e. when one wants reads that are 50%, 70% or 85% similar to the target (Figure 1). Relying only on input parameters such as gap penalties and seed size, which most well-known aligners have, it is difficult to achieve flexible mappings with require percentage of similarity. On the other hand, it is often desirable to find reads that are at least 90% similar to the provided target, and to discard the rest.

Another problem arises if there is insufficient data on the edge of the target located within a reference genome (Figure 1). In this situation the whole read (marked red) can potentially be discarded due to lack of data on the edge, even if a part of the read has significant similarity to the target.

We present SlopMap, bioinformatic software utility that quickly and flexibly identifies sequence reads that are within a given percent similarity to a target sequence. SlopMap is not a sequence alignment mapper, but rather identifies reads which may have been derived from

the target region. Unlike traditional alignment software, SlopMap only reports reads that are similar to the provided target. SlopMap selects reads for downstream analysis, such as assembly of sub-genome targets i.e. bacterial plasmids, virae, mitochondria, exome cap-ture data, chloroplasts, transcriptomes, etc. It employs exact kmer matching, which we call sloppy mapping, without conducting the computationally expensive alignment stage of traditional mappers. SlopMap can be directly embedded in biological data processing

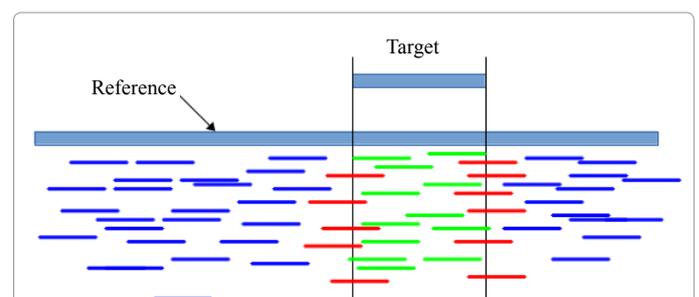


Figure 1: Situation when red reads on the edge of the target can be discarded by the general-purpose mapper. Green sequences can still be reported.

*Corresponding author: Ilya Y Zhbannikov, Institute for Bioinformatics and Evolutionary Studies, University of Idaho, USA, Tel: 208-885-6010; E-mail: zhba3458@vandals.uidaho.edu

Received May 07, 2013; Accepted June 24, 2013; Published July 02, 2013

Citation: Zhbannikov IY, Hunter SS, Settles ML, Foster JA (2013) SlopMap: A Software Application Tool for Quick and Flexible Identification of Similar Sequences Using Exact K-Mer Matching. J Data Mining Genomics Proteomics 4: 133. doi:10.4172/2153-0602.1000133

Copyright: © 2013 Zhbannikov IY, et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

pipelines before an assembly stage, since it maintains file format and preserves the original information such as bases, quality scores, and flowgrams (in the case of SFF files). SlopMap accepts both SFF (Roche 454 or Life Sciences Ion Torrent/Proton) and fastq (Illumina) file formats. SlopMap is a simple, easy to use and robust tool that can be used with percent similarity to targets as low as 5% (95% dissimilar). SlopMap along with its user manual is freely available under GPL from Bitbucket: <http://bitbucket.org/izhbannikov/slopmap/>.

Method

SlopMap is fully implemented and optimized in C++ for efficiency. This is a command line application with all the input parameters specified on the command line. SlopMap is tunable via input parameters for kmer size k , percentage of similarity t and distance between two consecutive kmers d . It also supports flexible input and output file formats: FastQ, FastA, SFF and TXT.

Input Files

Target: The target library file, also known as a “database”, is a FASTA formatted file that contains one or multiple records. Each record consists of two parts: a header and a sequence string. The header must contain a name which is a unique identifier of the record. The sequence string is DNA sequence which specifies the target of interest.

DNA Query Library Files: The DNA query library files are data from the NGS machines. SlopMap can take either Roche 454 SFF or FASTQ formatted files, or Illumina paired or single end reads. SlopMap computes the similarity of each read in a query library file to record in a target library.

Search procedure

SlopMap employs a multi-kmer search approach with single-base overlap (that can be increased) between two consecutive kmers in order to quickly determine the similarity to the target record. The algorithm is simple and straightforward:

1. Compute a dictionary from the given target library.
2. For each query sequence (“read”) in the DNA query library:
 - Compute a set of consecutive kmers.
 - Calculate the read’s similarity to the set of target sequences
3. Output summary statistics and files.

Dictionary: The first step is to build a dictionary indexed with kmers by sampling the given target (Figure 2). The kmers are short (usually 9-15 bases long) substrings representing the contiguous target. The target library sequences are sampled with the following pre-defined parameters: kmer size and a distance representing the constant overlap between consecutive kmers. These parameters remain unchanged throughout program execution. All kmers are hashed and associated with the offset position in the target string and a target record id. By default SlopMap uses kmer size of 15 bases. Google Dense Hash Map [9] that allows for fast data retrieval and memory efficiency is used as a data structure for the kmer dictionary.

Kmer matching: Query strings are sequentially sampled, so each query string contained within a given DNA query library is handled individually. Kmer size and a distance between two consecutive kmers remains the same for all reads in the query library. For each read taken from the query library, a dictionary search is performed and similarity between the read and the target is calculated as follows:

$$S = \frac{\text{number of shared kmers} * \text{kmer Length}}{\text{Min_length (query string, target string)}}, 0 \leq S \leq 1$$

Where number shared bases are the total length of kmers shared between the target and query string, Min length (query string, target string) is the minimum length, either of the query or target. Those reads that meet the pre-specified similarity threshold (parameter “ $-t$ ” in SlopMap, by default it is set to 0.75) are then saved in output files, others are ignored. The values of S range from 0, having no similar kmers to the target, to 1, having all kmers shared between the read and target.

Output files

The output files are:

- A report file that contains all information about the reads, including similarity, positions of first and last match within the target, bases and quality scores.
- One or two FASTQ formatted file(s), depending on whether the data are from single- or paired-end library.
- Optionally, an SFF formatted file, which contains only those Roche 454 sequences similar to the provided target.

Validation

To validate SlopMap we compared it to several alternative DNA mapping tools: Bowtie 2, BWA, Blat, on two different query DNA libraries: 621,578 Roche 454 *Escherichia coli K12 W583* reads; and 3,875,453 Illumina Yeast *Saccharomyces cerevisiae W303-K6001* reads. For the target sets, we randomly chose ten genes with various lengths from both genomes, each of which has over 4000 genes: thrA, thiQ, cydD, ycgB, dhaR, alka, ygfF, yphE, mscS, parC (E-coli, GenBank accession number: U00096.2); YNL095C, YNL094W, YNL093W, YNL092W, YNL091W, YNL090W, YNL089C, YNL088W, YNL087W, YNL085W (Yeast, GenBank accession number: AF458977.1).

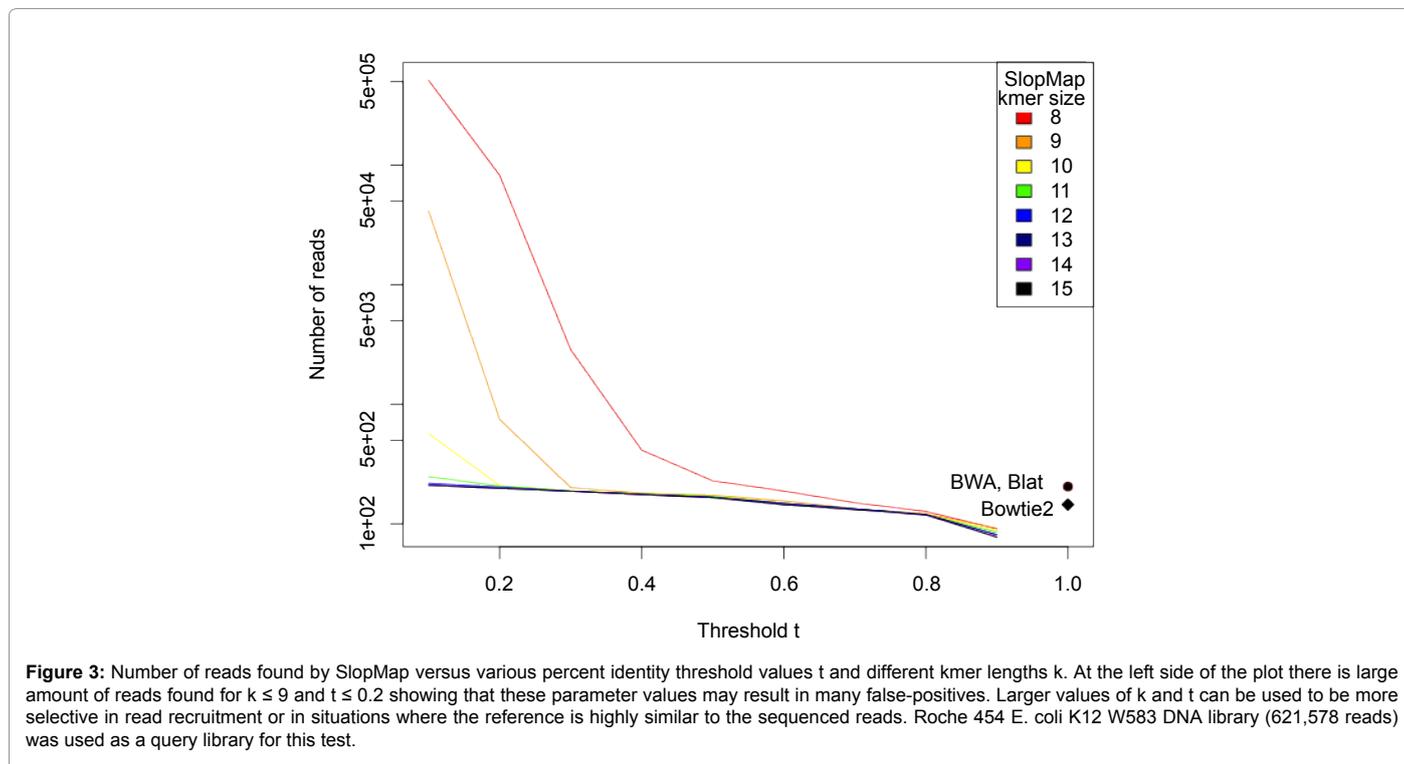
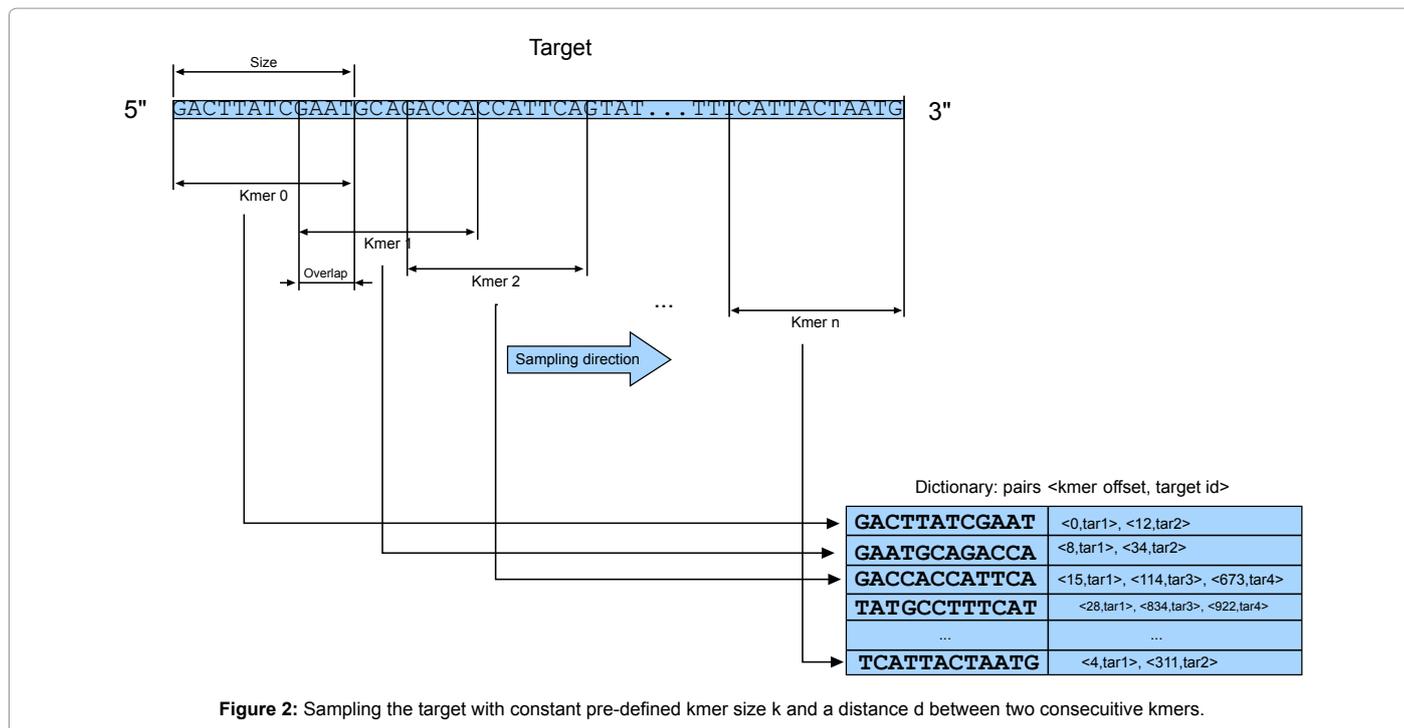
We estimate the number of reads found by SlopMap with various thresholds and kmers in order to:

- Estimate the effect of threshold to the number of reads found by SlopMap. Suggest values of t and k for optimal read recovery.
- Estimate the effect of various values of d (distance between two consecutive kmers) to the number of reads found by SlopMap against various threshold values allowing us to determine the range of optimal values for d .
- Compare the number of reads found by SlopMap to the number of reads found by other tools. In particular, to answer the question: what are the threshold values where read sets found by other tools are still subsets of reads found by SlopMap.

All tests were performed on a Linux server with Dual-Core AMD Opteron 8216 2.4 GHz processors (32 processors total) and 1 TB of shared memory and a laptop with single Intel Core i3 processor (four cores) with 4GB of memory.

Results and Discussion

We calculated the number of found reads using various threshold value and kmer sizes and compared our results to existing read mappers. These results are presented in Figure 3, which shows the number of reads against various distances between two consecutive kmers. Further comparisons are made to compare the overlap between recovered read sets. These results are presented as Venn diagrams in Figures 5 and 6. We also provided recommendation for optimal values of parameters k and d .



Number of reads found by SlopMap using different threshold values and kmer lengths

Figure 3 shows the number of reads found by SlopMap using different kmer sizes and threshold values. From this plot it's easy to see that for kmer values k , of 7, 8, 9 along with threshold value less than 0.1 (10% of similarity), large amount of reads found. These values of

k should be used only in such situations where the forged sequence is very divergent from the sequenced sample. Low values of t and k result in high sensitivity at the expense of specificity and should be used carefully to avoid multiple false positive hits.

Values of k within the range 10...15 can be used to generate more specific matches and are recommended for general usage. Higher

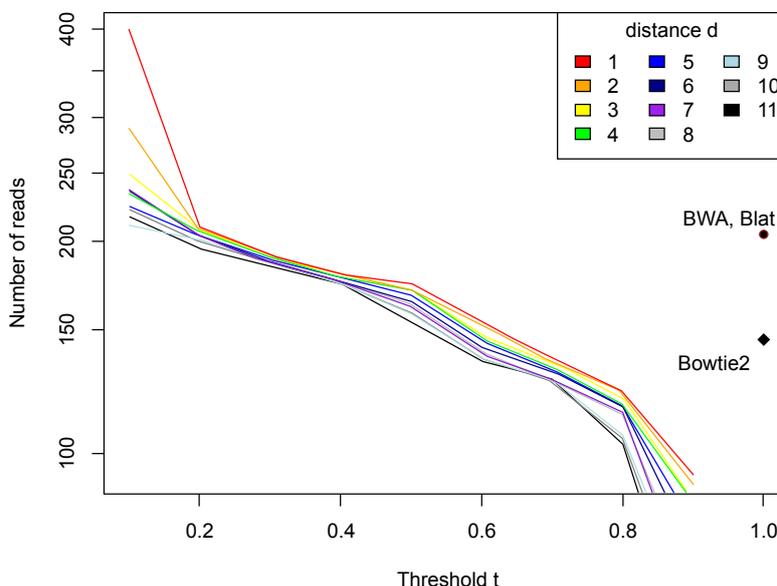


Figure 4: Number of reads found by SlopMap against various threshold values and different distances between two consecutive kmers. Roche 454 E. coli K12 W583 DNA library (621,578 reads) mapped against a single bacterial gene *ycgB* (1583 bp) was used as a query library for this test.

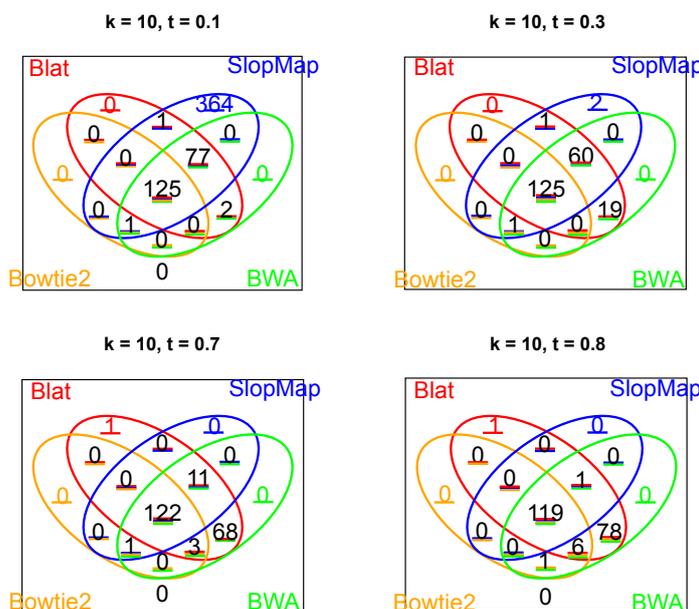


Figure 5: Number of reads reported by SlopMap and other tools depending on threshold values t and kmer sizes $k=10$ bases. When $t=0.3$ (the top-right figure), only 19 reads reported by Blat do not overlap with SlopMap, and those are also reported by BWA; reads reported by BWA and Bowtie2 are subsets of reads reported by SlopMap until $t \geq 0.7$. These tests were conducted on Roche 454 E. coli K12 W583 DNA library used as a query set.

values of k will result in less false positive mapping, SlopMap will not match kmers with mismatches, and will fail match reads at higher values of k . This is especially true in situations where reads have errors or adapters, which will generate false kmers and where there are real variants between the target and sequenced sample.

Number of reads found by SlopMap using various threshold values and distances between two consecutive kmers

In order to examine how different values of parameter d (the

distance between two consecutive kmers in a read) impacts to the number of reads found by SlopMap and propose the optimal value for d , we provided a set of tests with kmer size $k=11$ and threshold values $t=0.1...1.0$ (i.e. from 10% to 100% similarity). This is shown in Figure 4. We find that in this data set, d has minimal impact on read recruitment. However we observe a higher recruitment rate for lower d , suggesting higher sensitivity. Using a small d result in slower performance however, so in cases where target and reference are highly divergent, a low d should be used, while a higher d can be used for more similar

sequences. With these considerations in mind, we have set the default value of d to 3, and allow the user to change it using the command-line parameter $-d=N$. Our recommendations for the parameter d value to be no more than k mer length. Otherwise there may not be sufficient coverage. The values for the d from 1 to 5 are optimal for $k=11$ bases, since they give number of reads significantly higher than other tools within $t \leq 0.5$ ($t=0.75$ is set by default in SlopMap).

Sensitivity test: comparison the number of reads reported by SlopMap and other tools

Gaps occur when part of the query aligns to one part of the reference and another part aligns close to the first part but with a gap of one or more bases. Such gaps are usually well recognized by some widely used aligners. Another type of gap can occur at the end of a target sequence, when part of the query matches the target, resulting in an end gap. SlopMap can find such reads and thereby identify more similar sequences than some other alignment tools. In order to compare the sensitivity of SlopMap to a set of other mappers (BWA, Blat, Bowtie2), we conducted several tests using Roche 454 reads from *E. coli* mapped against *ycgB* gene sequence. We are interested what the cut-off point is when reads found by alignment tools are still subsets of reads found by SlopMap. We can roughly say that the set of reads reported by one application is a subset of reads reported by another application if there is more 95% overlap between these two sets. For two k mer lengths (10 and 15 bases) and threshold values (0.1...1 with step of 0.1) we computed Venn diagrams that show overlap sets reported by SlopMap and other tools. From these diagrams we conclude that for threshold values below 0.3 (30% similarity) and for both k mer sizes (10 and 15), the reads found by Blat is a subset of number of reads found by SlopMap. Threshold values when reads reported by Bowtie 2 and BWA are still subsets of number of reads reported by SlopMap are 0.7 and 0.3 (70% and 30% similarity) for k mer size 10 and 15 respectively. Results are shown in Figures 5 and 6. When threshold value t is 100%, SlopMap does not find any reads similar to the target sequence. This is expected

because reads can contain base-call errors, homopolymers and other artefacts that introduce noise into sequences.

Non-consecutive matches

When we compute the similarity of reads we do not assume that k mer matches are consecutive. Non-consecutive matching may occur in situation which is biologically possible such as exon shuffling, inversion, etc. In Figure 7, the read and a target are shown along with k mers shared between them. In this situation, k mer 1, 2 and 3 match corresponding k mers in a target but in different order (non-consecutive). The read can be still considered as similar to the target. SlopMap identifies and reports this read as being similar to the target, despite the rearrangement.

Timing considerations

Figure 8 displays the execution time required to complete each search. We compare execution times for various threshold values of SlopMap (other parameters were set to default) with other tools. SlopMap is faster than Bowtie2 and BWA, but slower than Blat, which is the fastest of the mappers we tested. However, Blat requires that the input be in FASTA format, and does not support writing output in FASTA or FASTQ format, making it necessary to perform additional steps, both before and after using the program. Post-mapping conversion work is also required for BWA and Bowtie2. Bowtie2 writes all sequences to the SAM file, and write unmapped or discordantly mapped reads to files using command line parameters to output mapped sequences in FASTQ file format.

Memory requirements

SlopMap is fast and require very little memory (2-200MB, depending on target size). The memory consumption of SlopMap during searching grows linearly with the number of sequences and also depends on the k mer size defined by the user. For example, when the *E. coli* data set containing sequences with a mean length of 450 characters

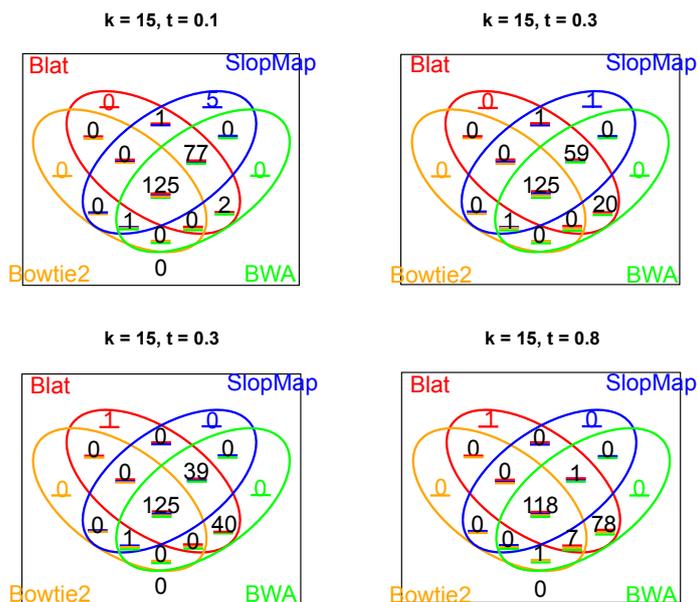


Figure 6: Number of reads reported by SlopMap and other tools depending on percent similarity threshold values t and k mer size $k = 15$ bases.

Non-consecutive k mer matching is still when using a target and query where rearrangements may have occurred. SlopMap identifies and reports these matches.

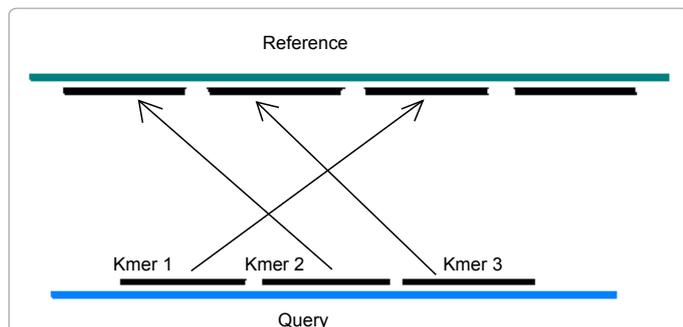


Figure 7: Non-consecutive kmer matching is still when using a target and query where rearrangements may have occurred. SlopMap identifies and reports these matches.

Search duration. Comparison of execution time between search tools. SlopMap is consistently fast across a range of threshold values (t), and perform somewhat faster than Bowtie 2, and significantly faster than BWA while using the same amount of CPU resources.

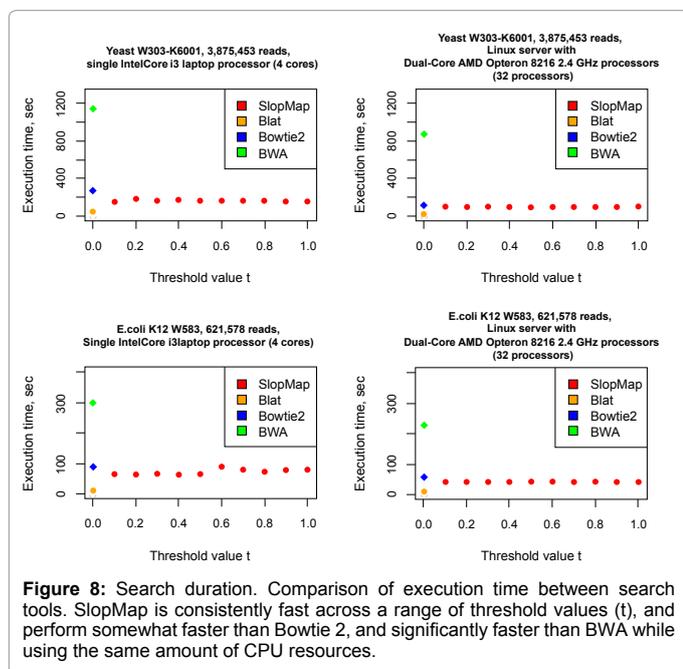


Figure 8: Search duration. Comparison of execution time between search tools. SlopMap is consistently fast across a range of threshold values (t), and perform somewhat faster than Bowtie 2, and significantly faster than BWA while using the same amount of CPU resources.

was indexed on 15-mers, 50 kB of memory was utilized for every 20,000 bases.

Conclusion and Future Work

SlopMap provides researchers with a high-throughput choice for searching large sets of reads against target sequences. The software presented is faster than some well-known aligners, sensitive to low-similarity matches when desired, and flexible enough to allow similarity comparison for DNA (and potentially RNA and proteins). SlopMap is specifically designed for matching queries against large (more than 500,000 sequences) query sets. Three of SlopMaps beneficial attributes are its speed, flexibility and ease of use. Despite being fast and efficient mapper, we plan to further improve SlopMap in the future by adding support for multicore execution and by exploring more space and time efficient methods for storing and looking up kmers.

We believe that the biological research community will benefit from using SlopMap.

Acknowledgments

This publication was made possible by NIH Grant P20 RR16454 from the INBRE Program of the National Center for Research Resources.

References

- Lulin H, Xiao Y, Pei S, Wen T, Shangqin H (2012) The first illumina-based de novo transcriptome sequencing and analysis of safflower flowers. *PLOS ONE*, 7: e38653.
- Guo Y, Li J, Li CI, Shyr Y, Samuels DC (2013) MitoSeek: extracting mitochondria information and performing high-throughput mitochondria sequencing analysis. *Bioinformatics*, 29: 1210-1211.
- Kent WJ (2002) BLAT--the BLAST-like alignment tool. *Genome Res* 12:656-664.
- Langmead B, Salzberg SL (2012) Fast gapped-read alignment with bowtie 2. *Nat Methods* 9: 357-359.
- Li H, Durbin R (2009) Fast and accurate short read alignment with burrows-wheeler transform. *Bioinformatics*, 25: 1754-1760.
- Li R, Yu C, Li Y, Lam TW, Yiu SM, et al. (2009) SOAP2: an improved ultrafast tool for short read alignment. *Bioinformatics*, 25: 1966-1967.
- The SAM Format Specification Working Group (2011) The sam format specification (v1.4-r985).
- <http://www.ncbi.nlm.nih.gov/Traces/trace.cgi?cmd=show&f=formats&m=doc&s=format#sff>.
- <https://code.google.com/p/sparsehash/>

This article was originally published in a special issue, **Bioinformatics for Highthroughput Sequencing** handled by Editor: Dr. Heinz Ulli Weier, Lawrence Berkeley National Laboratory, USA