# The Harvest suite for rapid core-genome alignment and visualization of thousands of intraspecific microbial genomes 

Todd J Treangen ${ }^{*+}$, Brian D Ondov ${ }^{\dagger}$, Sergey Koren and Adam M Phillippy


#### Abstract

Whole-genome sequences are now available for many microbial species and clades, however existing whole-genome alignment methods are limited in their ability to perform sequence comparisons of multiple sequences simultaneously. Here we present the Harvest suite of core-genome alignment and visualization tools for the rapid and simultaneous analysis of thousands of intraspecific microbial strains. Harvest includes Parsnp, a fast core-genome multi-aligner, and Gingr, a dynamic visual platform. Together they provide interactive core-genome alignments, variant calls, recombination detection, and phylogenetic trees. Using simulated and real data we demonstrate that our approach exhibits unrivaled speed while maintaining the accuracy of existing methods. The Harvest suite is open-source and freely available from: http://github.com/marbl/harvest.


## Rationale

Microbial genomes represent over $93 \%$ of past sequencing projects, with the current total over 10,000 and growing exponentially. Multiple clades of draft and complete genomes comprising hundreds of closely related strains are now available from public databases [1], largely due to an increase in sequencing-based outbreak studies [2]. The quality of future genomes is also set to improve as shortread assemblers mature [3] and long-read sequencing enables finishing at greatly reduced costs $[4,5]$.
One direct benefit of high-quality genomes is that they empower comparative genomic studies based on multiple genome alignment. Multiple genome alignment is a fundamental tool in genomics essential for tracking genome evolution [6-8], accurate inference of recombination [9-14], identification of genomic islands [15,16], analysis of mobile genetic elements [17,18], comprehensive classification of homology $[19,20]$, ancestral genome reconstruction [21], and phylogenomic analyses [22-24]. The task of whole-genome alignment is to create a catalog of relationships between the sequences of each genome (ortholog, paralog, xenolog, and so on [25]) to reveal their evolutionary history [26,27]. While several tools exist (LS-

[^0]BSR [28], Magic [29], Mavid [30], Mauve [31-33], MGA [34], M-GCAT [35], Mugsy [36], TBA [37], multi-LAGAN [38], PECAN [39]), multiple genome alignment remains a challenging task due to the prevalence of horizontal gene transfer [26,40], recombination, homoplasy, gene conversion, mobile genetic elements, pseudogenization, and convoluted orthology relationships [25]. In addition, the computational burden of multiple sequence alignment remains very high [41] despite recent progress [42].
The current influx of microbial sequencing data necessitates methods for large-scale comparative genomics and shifts the focus towards scalability. Current microbial genome alignment methods focus on all-versus-all progressive alignment $[31,36]$ to detect subset relationships (that is, gene gain/loss), but these methods are bounded at various steps by quadratic time complexity. This exponential growth in compute time prohibits comparisons involving thousands of genomes. Chan and Ragan [43] reiterated this point, emphasizing that current phylogenomic methods, such as multiple alignment, will not scale with the increasing number of genomes, and that 'alignment-free' or exact alignment methods must be used to analyze such datasets. However, such approaches do not come without compromising phylogenetic resolution [44].
Core-genome alignment is a subset of whole-genome alignment, focused on identifying the set of orthologous
sequence conserved in all aligned genomes. In contrast to the exponential complexity of multiple alignment, coregenome alignment is inherently more scalable because it ignores subset relationships. In addition, the core genome contains essential genes that are often vertically inherited and most likely to have the strongest signal-to-noise ratio for inferring phylogeny. The most reliable variants for building such phylogenies are single-nucleotide polymorphisms (SNPs). Thus, core-genome SNP typing is currently the standard method for reconstructing large phylogenies of closely related microbes [45]. Currently, there are three paradigms for core-genome SNP typing based on read mapping, k -mer analyses, and wholegenome alignment.

Read-based methods have dominated the bioinformatics methods landscape since the invention of high-fidelity, short-read sequencing ( 50 to 300 bp ) [46]. This has made it very affordable to sequence, yet extremely challenging to produce finished genomes [47,48]. Thus, comparative genomics has turned to highly efficient and accurate read mapping algorithms to carry out assembly-free analyses, spawning many mapping tools [49-52] and variant callers [53-55] for detecting SNPs and short Indels. Read-based variant calling typically utilizes a finished reference genome and a sensitive read mapper (BWA [51], Smalt), variant caller (samtools/bcftools [55], GATK [53]), and variant filter (minimum mapping quality, core genomic regions). This method has been shown effective in practice [56] and does not rely on assembly. However, mapping requires the read data, which is not always available and can be orders of magnitude larger than the genomes themselves. In addition, mapping can be sensitive to contaminant, overlook structural variation, misalign low-complexity and repetitive sequence, and introduce systematic bias in phylogenetic reconstruction [57-59].

Exact alignment methods, often formulated as k-mer matching, can produce high precision results in a fraction of the time required for gapped alignment methods [60-62]. Spectral k-mer approaches have been used to estimate genome similarity [63], and k-mer based methods are commonly used to identify or cluster homologous genomic sequence $[64,65]$. Recently, k-mers have also been extended to SNP identification. kSNP [66] identifies oddlength k -mers between multiple samples that match at all but the central position. The matched k-mers are then mapped back to a reference genome to locate putative SNPs. Conveniently, this approach is suitable for both assembled genomes and read sets, but sensitivity is sacrificed for the improved efficiency of exact alignment [67].
Genome assembly [4,68-75], followed by whole-genome alignment [38,76,77], is the original method for variant detection between closely related bacterial genomes [78] and has been shown to perform well across multiple sequencing platforms [79]. In addition to SNPs, whole-genome
alignment is able to reliably identify insertions and deletions (Indels) and other forms of structural variation. Thus, whole-genome alignment is the gold standard for comprehensive variant identification, but relies on highly accurate and continuous assemblies, which can be expensive to generate. Lastly, and unlike reference mapping, whole-genome alignment is not easily parallelized or scaled to many genomes.
Specifically for the task of whole-genome SNP typing, the choice of read- or genome-based methods can often depend on data availability. For example, of the 24,000 bacterial genomes currently in NCBI RefSeq [80], only $55 \%$ have associated SRA read data and analysis of the remaining $45 \%$ requires genome-based methods. Thankfully, recent advances in both sequencing technology and assembly algorithms are making microbial genomes more complete than ever before. Modern de Bruijn assemblers like SPAdes [81] are able to generate highquality assemblies from short reads [3], and long read technologies have enabled the automated finishing of microbial genomes for under $\$ 1,000$ [82]. With the number of publically available genomes currently doubling every 18 months [1], and genome quality improving with the arrival of new technologies, we set out to solve the problem of aligning thousands of closely-related whole genomes.

## Rapid core-genome alignment and visualization

Here we present Parsnp and Gingr for the construction and interactive visualization of massive core-genome alignments. For alignment, Parsnp combines the advantages of both whole-genome alignment and read mapping. Like whole-genome alignment, Parsnp accurately aligns microbial genomes to identify both structural and point variations, but like read mapping, Parsnp scales to thousands of closely related genomes. To achieve this scalability, Parsnp is based on a suffix graph data structure for the rapid identification of maximal unique matches (MUMs), which serve as a common foundation to many pairwise $[76,77,83]$ and multiple genome alignment tools [31-36]. Parsnp uses MUMs to both recruit similar genomes and anchor the multiple alignment. As input, Parsnp takes a directory of MultiFASTA files to be aligned; and as output, Parsnp produces a coregenome alignment, variant calls, and a SNP tree. These outputs can then be visually explored using Gingr. The details of Parsnp and Gingr are described below.

## MUMi recruitment

Parsnp is designed for intraspecific alignments and requires input genomes to be highly similar (for example, within the same subspecies group or $>=97 \%$ average nucleotide identity). For novel genomes or an inaccurate taxonomy, which genomes meet this criterion is not
always known. To automatically identify genomes suitable for alignment, Parsnp uses a recruitment strategy based on the MUMi distance [84]. Only genomes within a specified MUMi distance threshold are recruited into the full alignment.

## Compressed suffix graph

Parsnp utilizes a Directed Acyclic Graph (DAG) data structure, called a Compressed Suffix Graph (CSG), to index the reference genome for efficient identification of multi-MUMs. CSGs have the unique property of representing an optimally compressed structure, in terms of number of nodes and edges, while maintaining all intrinsic properties of a Suffix Tree. CSGs were originally proposed as a more space-efficient alternative to Suffix Trees and first implemented in M-GCAT [35]. Node and edge compression of the Suffix Tree incurs a lineartime construction penalty, but facilitates faster traversal of the structure once built. Provided sufficient memory, the CSG can be used to align genomes of any size; however, the current implementation has been optimized for microbial genomes, requiring approximately 32 bytes per reference base for CSG construction and 15 bytes per base for the aligned genomes. Note that because multi-MUMs are necessarily present in all genomes, the choice of a reference genome has no effect on the resulting alignment.

## Multi-MUM search

Once built for the reference genome, all additional genomes are streamed through the CSG, enabling rapid, linear-time identification of MUMs shared across all genomes. A divide-and-conquer algorithm, adapted from MGCAT [35], recursively searches for smaller matches and iteratively refines the multi-MUMs. Next, locally collinear blocks (LCBs) of multi-MUMs are identified. These LCBs form the basis of the core-genome alignment.

## Parallelized LCB alignment

The multi-MUMs within LCBs are used to anchor multiple alignments. Gaps between collinear multi-MUMs are aligned in parallel using MUSCLE [85]. To avoid the unnecessary overhead of reading and writing MultiFASTA alignment files, Parsnp makes direct library calls via a MUSCLE API. The MUSCLE library is packaged with Parsnp, but originally sourced from the Mauve code base [86]. As with Mauve, MUSCLE is used to compute an accurate gapped alignment between the match anchors. Though MUSCLE alignment can be computationally expensive, for highly similar genomes, the gaps between collinear multi-MUMs are typically very short (for example, a single SNP column in the degenerate case).

## SNP filtering and trees

The final Parsnp multiple alignment contains all SNP, Indel, and structural variation within the core genome. However, given their ubiquity in microbial genome analyses, Parsnp performs additional processing of the coregenome SNPs. First, all polymorphic columns in the multiple alignment are flagged to identify: (1) repetitive sequence; (2) small LCB size; (3) poor alignment quality; (4) poor base quality; and (5) possible recombination. Alignment quality is determined by a threshold of the number of SNPs and Indels contained within a given window size. Base quality is optionally determined using FreeBayes [54] to measure read support and mixed alleles. Bases likely to have undergone recent recombination are identified using PhiPack [87]. Only columns passing a set of filters based on these criteria are considered reliable core-genome SNPs. The final set of coregenome SNPs is given to FastTree2 [88] for reconstruction of the whole-genome phylogeny.

## Compressed alignment file

For simplicity and storage efficiency, the output of Parsnp includes a single binary file encoding the reference genome, annotations, alignment, variants, and tree. Thousand-fold compression of the alignment is achieved by storing only the columns that contain variants. The full multiple alignment can be faithfully reconstructed from this reference-compressed representation on demand. Since Parsnp focuses on aligning only core blocks of relatively similar genomes, the number of variant columns tends to increase at a sub-linear rate as the number of genomes increases, resulting in huge space savings versus alternative multiple alignment formats. Conversion utilities are provided for importing/exporting common formats to/from the binary archive file, including: BED, GenBank, FASTA, MAF, Newick, VCF, and XMFA.

## Interactive visualization

Developed in tandem with Parsnp, the visualization tool Gingr allows for interactive exploration of trees and alignments. In addition to the compressed alignment format, Gingr accepts standard alignment formats and can serve as a general-purpose multiple alignment viewer. Uniquely, Gingr is capable of providing dynamic exploration of alignments comprising thousands of genomes and millions of alignment columns. It is the first tool of its kind capable of dynamically visualizing multiple alignments of this scale. The alignment can be seamlessly zoomed from a display of variant density (at the genome level) to a full representation of the multiple alignment (at the nucleotide level). For exploration of phyletic patterns, the alignment is simultaneously presented along with the core-genome SNP tree, annotations, and dynamic variant highlighting. The tree can be zoomed by clade,
or individual genomes selected to expand via a fisheye zoom. Structural variation across the genome can also be displayed using Sybil coloring [89], where a color gradient represents the location and orientation of each LCB with respect to the reference. This is useful for identifying structurally variant regions of the core.

## Evaluation of performance

We evaluated Parsnp on three simulated datasets (derived from Escherichia coli (E. coli) K-12 W3110) and three real datasets (Streptococcus pneumoniae, Peptoclostridium difficile, and Mycobacterium tuberculosis). Parsnp is compared below versus two whole-genome alignment methods (Mugsy, Mauve), a k-mer based method (kSNP), and two commonly used mapping pipelines (based on Smalt and BWA). The Smalt pipeline replicates the methods of the landmark Harris et al. paper [90] that has been adopted in many subsequent studies. The BWA pipeline is similar to the Smalt pipeline, but uses BWA for read mapping (Materials and methods).

## Simulated E. coli W3110 dataset

To precisely measure the accuracy of multiple tools across varying levels of divergence, we computationally evolved the genome of E. coli K-12 W3110 at three different mutation rates: 0.00001 (low), 0.0001 (medium), and 0.001 (high) SNPs per site, per branch. An average of 10 rearrangements were introduced, per genome. Each dataset comprises 32 simulated genomes, forming a perfect binary tree. Approximately 65X coverage of Illumina MiSeq reads was simulated and assembled for each genome to create draft assemblies. For input, the whole-genome alignment programs were given the draft assemblies, and the mapping pipelines the raw reads. Additional file 1: Figure S1 details the computational performance on the simulated datasets. Parsnp was the only method to finish in fewer than 10 min on the 32genome dataset, with the other methods requiring between 30 min to 10 h . Table 1 gives the accuracy of each tool on each dataset. The tools were benchmarked using true-positive and false-positive rates compared to a known truth, which captures the full alignment accuracy. Figure 1 plots the performance of all tools averaged across all mutation rates.
The whole-genome alignment methods performed comparably across all three mutation rates (Figure 1, red squares), with Mauve exhibiting the highest sensitivity (97.42\%) and Parsnp the highest precision (99.99\%). In addition, Parsnp identified $98 \%$ the rearrangement breakpoints within $1 \%$ of the rearrangement length (+/$50 \mathrm{bp})$ with no breakpoint spans. Mugsy demonstrated slightly higher sensitivity than Parsnp but with lower precision. Mugsy's lower precision was traced to a single fumA paralog [91] misalignment that generated a high
number of false-positive SNPs. All genome alignment methods were affected by misalignment of repeats and missing or low-quality bases in the assembly.
Performance of the individual methods was also measured in terms of branch SNP and length error with respect to the true phylogeny (Figure 2). These errors closely followed the false-negative and false-positive rates of each method, with no distinguishable pattern or branch biases. On draft genomes, precise methods such as Parsnp yielded underestimates of branch lengths while more aggressive methods like Mugsy resulted in more overestimates (outliers not shown). The aggressive methods also showed more variance in performance across branches.

## Comparison on closed genomes

Mugsy, Mauve, and Parsnp all performed similarly on finished genomes (Figures 1 and 2, green squares), offering a significant boost in sensitivity over both draft assemblies and reference mapping. Mugsy, Mauve, and Parsnp all exhibited near perfect false-discovery rates (FDR), with Parsnp being the only method to not report a single false positive across the three datasets. Both Mauve and Mugsy were similarly near-perfect in terms of true-positive rates (TPR). The drop in sensitivity ( $0.9 \%$ ) for Parsnp on full genomes can be explained by a lack of an LCB extension method. Mugsy was the most affected by draft genomes, going from best on closed genomes to demonstrating more false positives (Table 1) and LCB counts (Table 2) on draft genomes. Parsnp offered the overall best FDR of the genome alignment methods, and the fewest number of LCBs, averaged across both draft and closed genome datasets.

## Comparison to read mapping methods

On average, mapping-based methods were as precise and $0.5 \%$ to $1 \%$ more sensitive than alignment of draft genomes (Figure 1, blue triangles). Smalt showed the highest sensitivity, while BWA was the most specific. The precision of the mapping approaches may be overestimated for this dataset due to the absence of noncore sequence that is known to confound mapping [58]. Parsnp was the only genome alignment method to match the precision of mapping, but with a slight reduction in sensitivity. However, when provided with finished genomes, the whole-genome alignment methods excel in both sensitivity and specificity compared to read mapping. Thus, the performance divide between wholegenome alignment and mapping is entirely due to assembly quality and completeness. Using short reads, both the mapping and assembly-based approaches suffer false negatives due to ambiguous mappings or collapsed repeats, respectively. Exceeding 99\% sensitivity for this test set requires either longer reads (for mapping) or

Table 1 Core-genome SNP accuracy for simulated E. coli datasets

| Method | Description $^{\text {a }}$ | FP <br> Low | FN <br> Low | FP <br> Med | FN <br> Med | FP <br> High | FN <br> High | TPR | FDR |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| Mauve | WGA | 148 | 318 | 198 | 2,877 | 100 | 30,378 | 0.974 | 0.0004 |
| Mauve (c) | WGA | 0 | 0 | 2 | 38 | 6 | 649 | 0.999 | 0 |
| Mugsy | WGA | $1,261^{\text {b }}$ | 395 | 1,928 | 3,371 | 1,335 | 34,923 | 0.970 | 0.0036 |
| Mugsy (c) | WGA | 2 | 0 | 2 | 0 | 1 | 81 | 0.999 | 0 |
| Parsnp | CGA | 23 | 423 | 45 | 3,494 | 7 | 35,466 | 0.970 | 0.0001 |
| Parsnp (c) | CGA | 0 | 24 | 0 | 603 | 0 | 10,989 | 0.992 | 0 |
| KSNP | KMER | 259 | 600 | 908 | 19,730 | 1,968 | 916,127 | 0.280 | 0.0086 |
| Smalt | MAP | 33 | 110 | 0 | 1,307 | 55 | 22,957 | 0.981 | 0.0001 |
| BWA | MAP | 0 | 168 | 16 | 1,947 | 27 | 27,091 | 0.9775 | 0.0000 |

Data shown indicates performance metrics of the evaluated methods on the three simulated E. coli datasets (low, medium, and high). Method: Tool used. (c) indicates aligner ran on closed genomes rather than draft assemblies.

False positive (FP) and false negative (FN) counts for the three mutation rates (low, med, and high). True positive rate TPR: TP/(TP + FN). False discovery rate FDR: FP/(TP + FP). A total of 1,299,178 SNPs were introduced into the 32-genome dataset, across all three mutational rates.
${ }^{\text {aparadigm }}$ employed by each method.
${ }^{\text {b }}$ Mugsy's lower precision was traced to a paralog misalignment that resulted in many false-positive SNPs.
CGA: core genome alignment, FN, number of truth SNP calls not detected, FP, number of SNP calls that are not in truth set, KMER: k-mer based SNP calls, MAP: read mapping, TP: number of SNP calls that agreed with the truth, WGA: whole-genome alignment.
complete genomes (for alignment) to accurately identify SNPs in the repetitive regions.

## Comparison on 31 Streptococcus pneumoniae genomes

Parsnp was compared to whole-genome alignment methods using the 31 -genome S. pneumoniae dataset
presented in the original Mugsy publication [36]. Angiuoli and Salzberg compared Mugsy, Mauve, and Nucmer + TBA to measure the number of LCBs and size of the core genome aligned. On this dataset, Parsnp aligns $90 \%$ of the bases aligned by Mugsy, while using 50\% fewer LCBs (Table 3). In addition, Parsnp ran


Figure 1 Core-genome SNP accuracy for simulated E. coli datasets. Results are averaged across low, medium, and high mutation rates. Red squares denote alignment-based SNP calls on draft assemblies, green squares alignment-based SNP calls on closed genomes, and blue triangles for read mapping. Full results for each dataset are given in Table 1.


Figure 2 Branch errors for simulated $E$. coli datasets. Simulated E. coli trees are shown for medium mutation rate ( 0.0001 per base per branch). (A) shows branch length errors as bars, with overestimates of branch length above each branch and underestimates below each branch. Maximum overestimate of branch length was $2.15 \%$ (bars above each branch) and maximum underestimate was $4.73 \%$ (bars below each branch). (B) shows branch SNP errors as bars, with false-positive errors above each branch and false-negative errors below each branch. The maximum FP SNP value is 6 (bars above each branch) and maximum FN SNP value is 23 (bars below each branch). Note that the bar heights have been normalized by the maximum value for each tree and are not comparable across trees. Outlier results from Mugsy were excluded from the branch length plot, and kSNP results are not shown. All genome alignment methods performed similarly on closed genomes, with Mauve and Mugsy exhibiting the best sensitivity (Table 1).
hundreds of times faster than the other methods, finishing this 31-way alignment in less than 60 s.

## Peptoclostridium difficile outbreak in the UK

Parsnp and Gingr are particularly suited for outbreak analyses of infectious diseases. To demonstrate this, we
applied Parsnp to a recent $P$. difficile outbreak dataset [92]. To generate input suitable for Parsnp, we assembled all genomes using iMetAMOS [93]. It is important to note that this was a resequencing project not intended for assembly and represents a worst case for a core-genome alignment approach; reads ranged from 50

Table 2 Comparison of locally collinear alignment block (LCB) count for simulated $E$. coli datasets, on assembled and finished genomes

| Method | Low | Medium | High |
| :--- | :--- | :--- | :--- |
| Mauve | 325 | 363 | 519 |
| Mauve (c) | 150 | 174 | 333 |
| Mugsy | 10,977 | 11,194 | 16,632 |
| Mugsy (c) | 237 | 247 | 351 |
| Parsnp | 205 | 271 | 344 |
| Parsnp (c) | 139 | 190 | 506 |

Method: Tool used.
(c) indicates aligner ran on closed genomes rather than draft assemblies. Low: $>99.99 \%$ similarity, Medium: $>99.9 \%$ similarity, High: $>99 \%$ similarity
to 100 bp in length and some genomes were sequenced without paired ends. The 826-way core genome alignment resulted in 1.4 Gbp being aligned in less than 5 h . The core genome represented $40 \%$ of the P. difficile 630 reference genome, consistent with previous findings [94]. Specifically, previous microarray experiments have indicated that $39 \%$ of the total CDS in the evaluated $P$. difficile clade pertains to the core genome ( $1 \%$ less than identified by Parsnp). Figure 3 shows a Gingr visualization of the 826 -way alignment and clade phylogeny. Related outbreak clusters are immediately visible from the phyletic patterns of the alignment, confirming the primary clades of the tree. In addition, the SNP heatmap highlights the phyletic signature of several subclades, in this case within the known hpdBCA operon [95] that is extremely well conserved across all 826 genomes.
Figure 4 shows a zoomed view of the 826 P. difficile genome alignment in Gingr, highlighting a single annotated gene. Although no metadata is publically available for this outbreak dataset, we identified that $b a c A$, a gene conferring antibiotic resistance to bacitracin, is conserved in all 826 isolates. While alternative antibiotic treatments for $P$. difficile infections have been wellstudied over the past 20 to 30 years [96], a recent study reported that $100 \%$ of 276 clinical isolates had high-level resistance to bacitracin [97]. In concordance with this study, our results indicate there may be widespread bacitracin resistance across this outbreak dataset. Thus

Table 3 Comparison to the 31 S. pneumoniae Mugsy benchmark

| Method | Time | Core (bp) | LCBs |
| :--- | :--- | :--- | :--- |
| Parsnp | 0.3 min | $1,428,407$ | 1,171 |
| Mugsy | 100 min | $1,590,820$ | 2,394 |
| Mauve | 377 min | $1,568,715$ | 1,366 |
| NUCmer + TBA | 80 min | $1,457,575$ | 27,075 |

Time: Method runtime from input to output. Core: Size of the aligned core genome measured in base pairs. LCBs: Number of locally colinear blocks in the alignment.
alternative antibiotics, such as vancomycin, could represent better treatment options.

## Mycobacterium tuberculosis geographic spread

For a second case evaluation, we ran Parsnp on a $M$. tuberculosis global diversity dataset [98]. In this case, the raw SNP calls were kindly made available (Iñaki Comas, personal communication), facilitating a direct comparison to the published results. The variant pipeline of Comas et al. is similar to our BWA pipeline, but with all SNP calls intersected with MAQ SNPfilter, which discards any SNP with neighboring Indels $\pm 3 \mathrm{bp}$ or surrounded by $>3 \mathrm{SNPs}$ within a 10 bp window. To replicate this study using whole-genome alignment, we assembled all genomes from the raw reads using iMetAMOS and ran Parsnp on the resulting draft assemblies. Figure 5 summarizes the results of the comparison and Figure 6 shows a Gingr visualization of the resulting tree and alignment, with major clades confirmed by correlations in the SNP density display.
Given a lack of truth for this dataset, we constructed a reference phylogeny based on the intersection of the Parsnp and Comas et al. SNP sets, which excludes potential false positives produced by only one of the methods. We evaluated the accuracy of phylogenetic reconstruction by measuring the Robinson-Foulds distance [99] and calculating the number of shared splits between the resulting trees (Figure 5). The Parsnp generated phylogeny has a higher percentage of shared splits with the reference phylogeny ( $94.1 \%$ versus $92.3 \%$ for Comas), while both methods exhibited a similar Robinson-Foulds distance to the reference phylogeny $(0.0007)$.
When compared directly, Parsnp was found to share 90,295 of its 104,896 SNPs calls ( $86 \%$ ) and 19,838 of its 22,131 SNP positions (90\%) with the Comas et al. study, resulting in $90.5 \%$ shared splits between the reconstructed phylogenies. We further investigated the SNP calls unique to either Parsnp or Comas and found that the majority formed a few well-defined splits that were largely consistent with the reference phylogeny (Additional file 1: Table S1). These SNPs are likely true positives that were missed by one of the two methods. A smaller fraction of unique SNPs formed single-genome splits, representing potential false positives due to sequencing or mapping error $(1,503$ for Parsnp, 3,016 for Comas).

## Runtime and storage analysis

To evaluate Parsnp's scalability, we profiled performance across six datasets ranging from 32 genomes to 10,000 genomes. Runtime was observed to increase linearly with additional genomes (Additional file 1: Figure S2), requiring a few minutes for the 32 genome E. coli dataset, 1.5 h for the 826 genome $P$. difficile dataset, and a


Figure 3 Gingr visualization of $\mathbf{8 2 6}$ P. difficile genomes aligned with Parsnp. The leaves of the reconstructed phylogenetic tree (left) are paired with their corresponding rows in the multi-alignment. A genome has been selected (rectangular aqua highlight), resulting in a fisheye zoom of several leaves and their rows. A SNP density plot (center) reveals the phylogenetic signature of several clades, in this case within the fully-aligned hpd operon (hpdB, hpdC, hpdA). The light gray regions flanking the operon indicate unaligned sequence. When fully zoomed (right), individual bases and SNPs can be inspected.
maximum of roughly 14 h to align the 10,000 genome set on a $2.2 \mathrm{GHz}, 32$-core, 1 TB RAM server (Table 4). In addition, for the 32 -genome simulated $E$. coli datasets, Parsnp was 10 to 100 times faster than all other methods evaluated. Maximum memory usage was 2 GB for the 145 Mbp E. coli dataset and 309 GB for the $21 \mathrm{Gbp} S$. pneumoniae dataset (Table 4). Memory usage can be explicitly limited via a command-line parameter (--max-partition-size) but this results in increased runtime.
In addition to runtime efficiency, Parsnp requires much less storage than the other approaches due to its binary alignment format and the compressive effect of assembly. For the 32-genome E. coli dataset, Parsnp's output totals just 4.5 MB , compared to 13 GB required to store compressed FASTQ [100] and VCF [101] files and 149 MB to store XMFA [38]. Storage reductions are amplified for larger datasets. For example, the raw read data for the $P$. difficile dataset requires 1.4 TB of storage (0.6 TB compressed). Assembling these data reduces the total to 3.3 GB by removing the redundancy of the reads. The XMFA alignment of these assemblies is 1.4 GB , and
reference-compressed binary format occupies just 15 MB . This equates to roughly a $100,000 \mathrm{X}$ (lossy) compression factor from raw reads to compressed archive, requiring only 0.08 bits per base to store the full core-genome alignment plus other related information, which is competitive with related techniques like CRAM [102]. As outbreak studies continue to expand in scale, whole-genome assembly and alignment presents a sustainable alternative to the current mapping-based strategies.

## Discussion

Parsnp is orders of magnitude faster than current methods for whole-genome alignment and SNP typing, but it is not without limitations. Parsnp represents a compromise between whole-genome alignment and read mapping. Compared to whole-genome aligners, Parsnp is less flexible because it is designed to conservatively align the core genome and is less sensitive as a result. Additionally, Parsnp is currently limited to intraspecific analysis of closely-related species and strains; future improvements will target more sensitive match detection for increased


Figure 4 Conserved presence of bacA antiobiotic resistance gene in P. difficile outbreak. Gingr visualization of conserved bacitracin resistance gene within the Parsnp alignment of 826 P. difficile genomes. Vertical lines indicate SNPs, providing visual support of subclades within this outbreak dataset.
tolerance of divergent genomes. Compared to read mapping, Parsnp is less robust and requires high-quality assemblies to maximize sensitivity. Thus, the right tool depends on the data and task at hand.

Core-genome alignment and phylogeny reconstruction are critical to microbial forensics and modern epidemiology. When finished or high-quality genomes are available, Parsnp is both efficient and accurate for these tasks. In addition, even for fragmented draft assemblies, Parsnp exhibits a favorable compromise between sensitivity and specificity. Surprisingly, Parsnp matched the specificity of the mapping-based approaches on the simulated datasets. However, multiplexed shortread sequencing followed by mapping still remains the most economical approach for sensitive analysis of large strain collections. Furthermore, whole-genome alignment depends on genome assemblies for variant detection; the accuracy of these methods for SNP calling is directly proportional to the quality of the assembly. Thus, Parsnp is recommended for analyzing high-quality assemblies or when raw read data are not available.
Assembled genomes have a number of advantages over read data - primarily compression and convenience. Storing, sharing, and analyzing raw read datasets incurs significant overhead from the redundancy in
sequencing (often 100-fold), and this burden nearly resulted in the closure of the NCBI SRA database [103]. Adding additional orders of magnitude to the already exponential growth of sequencing data is not sustainable. Instead, information in the reads not currently stored in common assembly formats (for example, allelic variants) should be propagated to the assembled representation, forming a compressed, but nearly lossless format. In this way, genomes could be shared in their native, assembled format, saving both space and time of analysis. Here, we have taken a small step in that direction by identifying low quality bases, as computed by FreeBayes [54]. This allows filtering of low quality and mixed alleles and improves the specificity of the assembly-based approaches. However, more comprehensive, graph-based formats are needed to capture the full population information contained in the raw reads.

Parsnp was also built around the observation that high-quality, finished genome sequences have become more common as sequencing technology and assembly algorithms continue to improve. New technologies, such as PacBio SMRT sequencing [104] are enabling the generation of reference-grade sequences at extremely reduced costs. This presents another opportunity for Parsnp - the construction and maintenance of core


Figure 5 Comparison of Parsnp and Comas et al. result on M. tuberculosis dataset. A Venn diagram displays SNPs unique to Comas et al. [98] (left, blue), unique to Parsnp (right, red), and shared between the two analyses (middle, brown). On top, an unrooted reference phylogeny is given based on the intersection of shared SNPs produced by both methods ( $90,295 \mathrm{SNPs}$ ). On bottom, the phylogenies of Comas et al. (left) and Parsnp (right) are given. Pairs of trees are annotated with their Robinson-Foulds distance (RFD) and percentage of shared splits. The Comas et al. and Parsnp trees are largely concordant with each other and the reference phylogeny. All major clades are shared and well supported by all three trees.
genomes and trees for clinically important species. With well-defined reference cores, outbreaks could be accurately typed in real-time by mapping sequences directly to the tree using phylogenetically aware methods such as pplacer [105] or PAGAN [106]. Such a phylogenetic approach would be preferable to alternative typing schemes based on loosely defined notions of similarity, such as pulse-field electrophoresis (PFGE) [107] and multi-locus sequence typing (MLST) [108].

## Conclusion

Parsnp offers a highly efficient method for aligning the core genome of thousands of closely related species, and Gingr provides a flexible, interactive visualization tool for the exploration of huge trees and alignments. Together, they enable analyses not previously possible with whole-genome aligners. We have demonstrated that Parsnp provides highly specific variant calls, even for highly fragmented draft genomes, and can efficiently


Figure 6 Gingr visualization of 171 M. tuberculosis genomes aligned with Parsnp. The visual layout is the same as Figure 3, but unlike Figure 3, a SNP density plot across the entire genome is displayed. Major clades are visible as correlated SNP densities across the length of the genome.
reconstruct recent outbreak analyses including hundreds of whole genomes. Future improvements in genome assembly quality and formats will enable comprehensive cataloging of microbial population variation, including both point and structural mutations, using genome alignment methods such as Parsnp.

## Materials and methods

## Software and configurations

Mugsy [36] v1.23 and Mauve Aligner [31,33] v2.3.1 were run using default parameters on assembled sequences. mauveAligner was selected instead of progressiveMauve due to improved performance on the simulated $E$. coli
datasets, which do not contain subset relationships. kSNP v2.0 [66] was run with a k-mer size of 25 on both the raw read data and the assemblies; the assemblies were merged with Ns using the provided merge_fasta_contigs.pl utility. Raw MAF/XMFA/VCF output was parsed to recover SNPs and build MultiFASTA files.
Smalt version 0.7 .5 was run with default parameters for paired reads, mirroring the pipeline used in several recent SNP typing studies [90,109-111]. Samtools view was used to filter for alignments with mapping qualities greater than or equal to 30 . Variants were called by piping samtools mpileup output into bcftools view with the -v (variants only), -g (genotype) and -I (skip

Table 4 Performance profile of Parsnp runtime (MUM + alignment) on all evaluated datasets

| Dataset | Num Genomes ${ }^{\text {a }}$ | Aligned ${ }^{\text {b }}$ (Mbp) | MUM ${ }^{\text {c }}$ (min) | MUSCLE ${ }^{\text {d }}$ (min) | Total ${ }^{\text {e }}$ (min) | Mem ${ }^{\text {f }}$ (GB) |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| E. coli (avg) | 32 | 142 | 2 | 2 | 4 | 2 |
| M. tuberculosis | 171 | 424 | 12 | 20 | 32 | 14 |
| P. difficile | 826 | 1,392 | 46 | 39 | 85 | 71 |
| S. aureus SIM | 10,000 | 21,000 | 668 | 201 | 869 | 309 |

Results were generated on a 32-core, $2.2 \mathrm{GHz}, 1$ TB RAM Linux server. Dataset: the genome set.
${ }^{\text {a }}$ The number of genomes aligned.
${ }^{5}$ Total Mbp aligned.
${ }^{\text {'The }}$ The time spent finding maximal unique matches.
${ }^{\mathrm{d}}$ The time spent performing gapped multi-alignment with MUSCLE.
${ }^{\text {e }}$ Total Parsnp runtime (sum of MUM and MUSCLE).
${ }^{\mathrm{f}}$ Maximum memory usage.

Indels) flags. Variants were then filtered with VCFUtils varFilter with the -d (minimum read depth) parameter set to 3 . Variants for all samples of each set were called concomitantly by providing samtools mpileup with all BAM files.
BWA [52] was run in its standard paired-end alignment mode with default parameters, using aln to align each set of ends and sampe to produce a combined SAM file. Samtools view was used to filter for alignments with mapping qualities greater than or equal to 30 . Variants were called by piping samtools mpileup output into bcftools view with the -v (variants only), -g (genotype) and -I (skip Indels) flags. Variants were then filtered with VCFUtils varFilter with the -d (minimum read depth) parameter set to 3 . As with Smalt, variants for all samples of each set were called concomitantly by providing samtools mpileup with all BAM files.
FastTree v2 [88] was used to reconstruct phylogenies using default parameters.

## E. coli K-12 W3110 simulated dataset

The complete genome of E. coli K-12 W3110 [112], was downloaded from RefSeq (AC_000091). This genome was used as the ancestral genome and evolution was simulated along a balanced tree for three evolutionary rates using the Seq-Gen package [113] with parameters mHKY -t4.0-14646332-n1-k1 and providing the corresponding binary tree evolved at three evolutionary rates: $0.00001,0.0001$, and 0.001 SNPs per site, per branch. This corresponds to a minimum percent identity of approximately $99 \%, 99.9 \%$, and $99.99 \%$ between the two most divergent genomes, respectively, reflecting the variation seen in typical outbreak analyses. No small ( $<5 \mathrm{bp}$ ) or large Indels were introduced, but an average of 101 Kbp rearrangements (inversions and translocations) were added, per genome, using a custom script [114]. Paired reads were simulated to model current MiSeq lengths ( $2 \times 150 \mathrm{bp}$ ) and error rates (1\%). Moderate coverage, two million PE reads (64X coverage), was simulated for each of the 32 samples using wgsim (default parameters, no Indels), from samtools package version 0.1.17 [55].
Two of the simulated read sets were independently run through iMetAMOS [93] to automatically determine the best assembler. The consensus pick across both datasets was SPAdes version 3.0 [81], which was subsequently run on the remaining 30 simulated read sets using default parameters. The final contigs and scaffolds files were used as input to the genome alignment methods. For mapping methods, the raw simulated reads were used. For accuracy comparisons, Indels were ignored and called SNPs were required to be unambiguously aligned across all 32 genomes (that is, not part of a subset relationship; SNPs present but part of a subset relationship were ignored).

## S. pneumoniae dataset

A full listing of accession numbers for the 31-genome $S$. pneumoniae dataset is described in [36]. For scalability testing, Streptococcus pneumoniae TIGR4 (NC_003028.3) was used to create a pseudo-outbreak clade involving 10,000 genomes evolved along a star phylogeny with on average 10 SNPs per genome.

## M. tuberculosis dataset

We downloaded and assembled sequencing data from a recently published study of $M$. tuberculosis [98]. A total of 225 runs corresponding to project ERP001731 were downloaded from NCBI SRA and assembled using the iMetAMOS ensemble of SPAdes, MaSuRCA, and Velvet. The iMetAMOS assembly for each sample can be replicated with the following commands, which will automatically download the data for RUN_ID directly from SRA:

```
> initPipeline -d asmTB -W iMetAMOS -m RUN_ID -i
200:800
> runPipeline -d asmTB -a spades,masurca,velvet -p 16
```

The M. tuberculosis dataset included a mix of single and paired-end runs with a sequence length in the range of 51 to 108 bp . The average k-mer size selected for unpaired data was 26 , resulting in an average of 660 contigs and an N50 size of 17 Kbp . For paired-end data, the average selected k -mer was 35 , resulting in an average of 333 contigs and an N50 size of 43 Kbp . Assemblies containing more than 2,000 contigs, or 1.5 X larger/smaller than the reference genome, were removed. The final dataset was reduced to 171 genomes, limited to labeled strains that could be confidently matched to the strains used in the Comas et al. study for SNP and phylogenetic comparison.

## P. difficile dataset

Note, Clostridium difficile was recently renamed to Peptoclostridium difficile [115]. We downloaded and assembled sequencing data from a recently published study of P. difficile [92]. A total of 825 runs corresponding to project ERP003850 were downloaded from NCBI SRA [86] and assembled within iMetAMOS this time only using SPAdes, which was identified as the best performer on the $M$. tuberculosis dataset. The iMetAMOS assembly for each sample can be replicated with the following commands, which will download the data for RUN_ID directly from SRA:

```
> initPipeline -d asmPD -W iMetAMOS -m RUN_ID -i
200:800
> runPipeline -d asmPD -a spades -p 16
```

The $P$. difficile dataset included paired-end runs with a sequence length in the range of 51 to 100 bp . SPAdes was
selected as the assembler and run with k -mer sizes of 21 , 33,55 , and 77 . The assemblies had an average of 660 contigs and an N50 size of 138 Kbp . Assemblies containing more than 2,000 contigs, or 1.5 X larger/smaller than the reference genome, were removed.

## Data and software availability

All data, supplementary files, assemblies, packaged software binaries and scripts described in the manuscript are available from: https://www.cbcb.umd.edu/software/ harvest. The python script used to introduce rearrangements into the simulated genomes is also available for download at: https://github.com/marbl/parsnp/tree/master/script. Source code of the described software, including Parsnp and Gingr, is available for download from: http://github.com/marbl/harvest.

## Additional file

Additional file 1: Table S1. SNPs unique to each method characterized by the five most common splits found. Figure S1. Runtime comparison for the whole-genome alignment methods on the simulated 32-genome E. coli W3110 dataset. Figure S2. Timing performance from 32 to 10,000 S. pneumoniae genomes.

## Abbreviations

Bp: base pair; ERA: European Read Archive; Indel: insertion or deletion; LCB: locally collinear block; Mbp: million base pairs; MUM: maximal unique match; MUMi: similarity index based on maximal unique matches; NGS: Next generation sequencing; PE: paired-end; SNP: single-nucleotide polymorphism; SRA: Sequence Read Archive; VCF: variant call format; XMFA: extendend multi-fasta format.

## Competing interests

The authors declare that they have no competing interests.

## Authors' contributions

TJT, BDO, and AMP conceived the method, designed the experiments, and drafted the manuscript. TJT implemented Parsnp and BDO implemented Gingr. TJT, BDO, and SK performed the experiments. All authors read and approved the final manuscript

## Acknowledgments

We would like to thank Iñaki Comas for providing the raw data for the TB study; Xiangyu Deng for helpful comments on the manuscript and software; and Bradd Haley, Bill Klimke, Jason Sahl, Maliha Aziz, and Suzanne Bialek-Davenet for feedback on early versions of the software. This work was funded under Agreement No. HSHQDC-07-C-00020 awarded by the Department of Homeland Security Science and Technology Directorate (DHS/S\&T) for the management and operation of the National Biodefense Analysis and Countermeasures Center (NBACC), a Federally Funded Research and Development Center. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security. In no event shall the DHS, NBACC, or Battelle National Biodefense Institute (BNBI) have any responsibility or liability for any use, misuse, inability to use, or reliance upon the information contained herein. The Department of Homeland Security does not endorse any products or commercial services mentioned in this publication

Received: 24 September 2014 Accepted: 31 October 2014
Published online: 19 November 2014

## References

1. Pagani I, Liolios K, Jansson J, Chen IM, Smirnova T, Nosrat B, Markowitz VM Kyrpides NC: The Genomes OnLine Database (GOLD) v. 4: status of genomic and metagenomic projects and their associated metadata. Nucleic Acids Res 2012, 40:D571-D579
2. Rasko DA, Webster DR, Sahl JW, Bashir A, Boisen N, Scheutz F, Paxinos EE, Sebra R, Chin CS, Iliopoulos D, Klammer A, Peluso P, Lee L, Kislyuk AO, Bullard J, Kasarskis A, Wang S, Eid J, Rank D, Redman JC, Steyert SR, FrimodtMoller J, Struve C, Petersen AM, Krogfelt KA, Nataro JP, Schadt EE, Waldor MK: Origins of the E. coli strain causing an outbreak of hemolytic-uremic syndrome in Germany. N Engl J Med 2011, 365:709-717.
3. Magoc T, Pabinger S, Canzar S, Liu X, Su Q, Puiu D, Tallon LJ, Salzberg SL: GAGE-B: an evaluation of genome assemblers for bacterial organisms. Bioinformatics 2013, 29:1718-1725.
4. Koren S, Schatz MC, Walenz BP, Martin J, Howard JT, Ganapathy G, Wang Z, Rasko DA, McCombie WR, Jarvis ED, Adam MP: Hybrid error correction and de novo assembly of single-molecule sequencing reads. Nat Biotechnol 2012, 30:693-700.
5. Chin CS, Alexander DH, Marks P, Klammer AA, Drake J, Heiner C, Clum A Copeland A, Huddleston J, Eichler EE, Turner SW, Korlach J: Nonhybrid, finished microbial genome assemblies from long-read SMRT sequencing data. Nat Methods 2013, 10:563-569.
6. Snitkin ES, Zelazny AM, Thomas PJ, Stock F, Group NCSP, Henderson DK, Palmore TN, Segre JA: Tracking a hospital outbreak of carbapenemresistant Klebsiella pneumoniae with whole-genome sequencing. Sci Transl Med 2012, 4:148ra116.
7. Gonzalez-Candelas F, Bracho MA, Wrobel B, Moya A: Molecular evolution in court: analysis of a large hepatitis $C$ virus outbreak from an evolving source. BMC Biol 2013, 11:76.
8. Kennemann L, Didelot X, Aebischer T, Kuhn S, Drescher B, Droege M, Reinhardt R, Correa P, Meyer TF, Josenhans C, Falush D, Suerbaum S: Helicobacter pylori genome evolution during human infection. Proc Natl Acad Sci U S A 2011, 108:5033-5038.
9. Yahara K, Didelot X, Ansari MA, Sheppard SK, Falush D: Efficient inference of recombination hot regions in bacterial genomes. Mol Biol Evol 2014, 31:1593-1605.
10. Namouchi A, Didelot X, Schock U, Gicquel B, Rocha EP: After the bottleneck: Genome-wide diversification of the Mycobacterium tuberculosis complex by mutation, recombination, and natural selection. Genome Res 2012, 22:721-734.
11. Didelot X, Meric G, Falush D, Darling AE: Impact of homologous and non-homologous recombination in the genomic evolution of Escherichia coli. BMC Genomics 2012, 13:256.
12. Didelot X, Bowden R, Street T, Golubchik T, Spencer C, McVean G, Sangal V, Anjum MF, Achtman M, Falush D, Donnelly P: Recombination and population structure in Salmonella enterica. PLoS Genet 2011, 7:e1002191.
13. Didelot X, Lawson D, Darling A, Falush D: Inference of homologous recombination in bacteria using whole-genome sequences. Genetics 2010, 186:1435-1449.
14. Didelot X, Darling A, Falush D: Inferring genomic flux in bacteria. Genome Res 2009, 19:306-317.
15. Dhillon BK, Chiu TA, Laird MR, Langille MG, Brinkman FS: IslandViewer update: Improved genomic island discovery and visualization. Nucleic Acids Res 2013, 41:W129-W132
16. Langille MG, Hsiao WW, Brinkman FS: Evaluation of genomic island predictors using a comparative genomics approach. BMC Bioinformatics 2008, 9:329.
17. Touchon M, Rocha EP: Causes of insertion sequences abundance in prokaryotic genomes. Mol Biol Evol 2007, 24:969-981.
18. Fuxelius HH, Darby A, Min CK, Cho NH, Andersson SG: The genomic and metabolic diversity of Rickettsia. Res Microbiol 2007, 158:745-753.
19. Dessimoz C, Gabaldon T, Roos DS, Sonnhammer EL, Herrero J: Quest for Orthologs C: Toward community standards in the quest for orthologs. Bioinformatics 2012, 28:900-904.
20. Gabaldon T, Dessimoz C, Huxley-Jones J, Vilella AJ, Sonnhammer EL, Lewis S: Joining forces in the quest for orthologs. Genome Biol 2009, 10:403.
21. Alekseyev MA, Pevzner PA: Breakpoint graphs and ancestral genome reconstructions. Genome Res 2009, 19:943-957.
22. Philippe H, Brinkmann H, Lavrov DV, Littlewood DT, Manuel M, Worheide G, Baurain D: Resolving difficult phylogenetic questions: why more sequences are not enough. PLoS Biol 2011, 9:e1000602.
23. Philippe H, Delsuc F, Brinkmann H, Lartillot N: Phylogenomics. Annu Rev Ecol Evol Syst 2005, 36:541-562.
24. Hayward A, Grabherr M, Jern P: Broad-scale phylogenomics provides insights into retrovirus-host evolution. Proc Natl Acad Sci U S A 2013, 110:20146-20151.
25. Gabaldon T, Koonin EV: Functional and evolutionary implications of gene orthology. Nat Rev Genet 2013, 14:360-366.
26. Dagan T, Artzy-Randrup Y, Martin W: Modular networks and cumulative impact of lateral transfer in prokaryote genome evolution. Proc Natl Acad Sci U S A 2008, 105:10039-10044.
27. Dale J, Price EP, Hornstra H, Busch JD, Mayo M, Godoy D, Wuthiekanun V, Baker A, Foster JT, Wagner DM, Tuanyok A, Warner J, Spratt BG, Peacock SJ, Currie BJ, Keim P, Pearson T: Epidemiological tracking and population assignment of the non-clonal bacterium. Burkholderia pseudomallei. PLoS Negl Trop Dis 2011, 5:e1381.
28. Sahl JW, Caporaso JG, Rasko DA, Keim P: The large-scale blast score ratio (LS-BSR) pipeline: a method to rapidly compare genetic content between bacterial genomes. Peer J 2014, 2:e332.
29. Swidan F, Rocha EP, Shmoish M, Pinter RY: An integrative method for accurate comparative genome mapping. PLoS Comput Biol 2006, 2:e75.
30. Bray N, Pachter L: MAVID: constrained ancestral alignment of multiple sequences. Genome Res 2004, 14:693-699.
31. Darling AE, Mau B, Perna NT: progressiveMauve: multiple genome alignment with gene gain, loss and rearrangement. PLoS One 2010, 5: 111147.
32. Darling AE, Treangen TJ, Messeguer X, Perna NT: Analyzing patterns of microbial evolution using the mauve genome alignment system. Methods Mol Biol 2007, 396:135-152.
33. Darling AC, Mau B, Blattner FR, Perna NT: Mauve: multiple alignment of conserved genomic sequence with rearrangements. Genome Res 2004, 14:1394-1403.
34. Hohl M, Kurtz S, Ohlebusch E: Efficient multiple genome alignment. Bioinformatics 2002, 18:S312-S320.
35. Treangen TJ, Messeguer X: M-GCAT: interactively and efficiently constructing large-scale multiple genome comparison frameworks in closely related species. BMC Bioinformatics 2006, 7:433.
36. Angiuoli SV, Salzberg SL: Mugsy: fast multiple alignment of closely related whole genomes. Bioinformatics 2011, 27:334-342.
37. Blanchette M, Kent WJ, Riemer C, Elnitski L, Smit AF, Roskin KM, Baertsch R, Rosenbloom K, Clawson H, Green ED, Haussler D, Miller W: Aligning multiple genomic sequences with the threaded blockset aligner. Genome Res 2004, 14:708-715.
38. Brudno M, Do CB, Cooper GM, Kim MF, Davydov E, Program NCS, Green ED, Sidow A, Batzoglou S: LAGAN and Multi-LAGAN: efficient tools for largescale multiple alignment of genomic DNA. Genome Res 2003, 13:721-731.
39. Paten B, Herrero J, Beal K, Fitzgerald S, Birney E: Enredo and Pecan: genome-wide mammalian consistency-based multiple alignment with paralogs. Genome Res 2008, 18:1814-1828.
40. Ochman H, Lerat E, Daubin V: Examining bacterial species under the specter of gene transfer and exchange. Proc Natl Acad Sci U S A 2005, 102:6595-6599.
41. Wang L, Jiang T: On the complexity of multiple sequence alignment. J Comput Biol 1994, 1:337-348.
42. Liu K, Warnow T: Large-scale multiple sequence alignment and tree estimation using SATe. Methods Mol Biol 2014, 1079:219-244.
43. Chan CX, Ragan MA: Next-generation phylogenomics. Biol Direct 2013, 8:3.
44. Posada D: Phylogenetic models of molecular evolution: next-generation data, fit, and performance. J Mol Evol 2013, 76:351-352.
45. Robinson ER, Walker TM, Pallen MJ: Genomics and outbreak investigation: from sequence to consequence. Genome Med 2013, 5:36.
46. Bennett S: Solexa Ltd. Pharmacogenomics 2004, 5:433-438.
47. Schatz MC, Delcher AL, Salzberg SL: Assembly of large genomes using second-generation sequencing. Genome Res 2010, 20:1165-1173.
48. Earl D, Bradnam K, St John J, Darling A, Lin D, Fass J, Yu HO, Buffalo V, Zerbino DR, Diekhans M, Nguyen N, Ariyaratne PN, Sung WK, Ning Z, Haimel M, Simpson JT, Fonseca NA, Birol I, Docking TR, Ho IY, Rokhsar DS, Chikhi R, Lavenier D, Chapuis G, Naquin D, Maillet N, Schatz MC, Kelley DR, Phillippy AM, Koren S, et al: Assemblathon 1: a competitive assessment of de novo short read assembly methods. Genome Res 2011, 21:2224-2241.
49. Langmead B: Aligning short sequencing reads with Bowtie. Curr Protoc Bioinformatics 2010, Chapter 11:Unit 1117.
50. Langmead B, Salzberg SL: Fast gapped-read alignment with Bowtie 2. Nat Methods 2012, 9:357-359.
51. Li H, Durbin R: Fast and accurate long-read alignment with BurrowsWheeler transform. Bioinformatics 2010, 26:589-595.
52. Li H, Durbin R: Fast and accurate short read alignment with BurrowsWheeler transform. Bioinformatics 2009, 25:1754-1760.
53. McKenna A, Hanna M, Banks E, Sivachenko A, Cibulskis K, Kernytsky A, Garimella K, Altshuler D, Gabriel S, Daly M, DePristo MA: The Genome Analysis Toolkit: a MapReduce framework for analyzing next-generation DNA sequencing data. Genome Res 2010, 20:1297-1303.
54. Garrison E, Marth G: Haplotype-based variant detection from short-read sequencing. arXiv 2012, 1207:3907. [http://arxiv.org/abs/1207.3907]
55. Li H, Handsaker B, Wysoker A, Fennell T, Ruan J, Homer N, Marth G, Abecasis G, Durbin R, Genome Project Data Processing S: The Sequence Alignment/ Map format and SAMtools. Bioinformatics 2009, 25:2078-2079.
56. Harris SR, Torok ME, Cartwright EJ, Quail MA, Peacock SJ, Parkhill J: Read and assembly metrics inconsequential for clinical utility of whole-genome sequencing in mapping outbreaks. Nat Biotechnol 2013, 31:592-594.
57. Bertels F, Silander OK, Pachkov M, Rainey PB, van Nimwegen E: Automated reconstruction of whole-genome phylogenies from short-sequence reads. Mol Biol Evol 2014, 31:1077-1088.
58. Li H : Toward better understanding of artifacts in variant calling from high-coverage samples. Bioinformatics 2014, 30:2843-2851.
59. Treangen TJ, Salzberg SL: Repetitive DNA and next-generation sequencing: computational challenges and solutions. Nat Rev Genet 2012, 13:36-46.
60. Vinga S, Almeida J: Alignment-free sequence comparison-a review. Bioinformatics 2003, 19:513-523.
61. Patro R, Mount SM, Kingsford C: Sailfish enables alignment-free isoform quantification from RNA-seq reads using lightweight algorithms. Nat Biotechnol 2014, 32:462-464.
62. Wood DE, Salzberg SL: Kraken: ultrafast metagenomic sequence classification using exact alignments. Genome Biol 2014, 15:R46.
63. Chor B, Horn D, Goldman N, Levy Y, Massingham T: Genomic DNA k-mer spectra: models and modalities. Genome Biol 2009, 10:R108.
64. Hauser M, Mayer CE, Soding J: kClust: fast and sensitive clustering of large protein sequence databases. BMC Bioinformatics 2013, 14:248.
65. Ghodsi M, Liu B, Pop M: DNACLUST: accurate and efficient clustering of phylogenetic marker genes. BMC Bioinformatics 2011, 12:271.
66. Gardner SN, Hall BG: When whole-genome alignments just won't work: kSNP v2 software for alignment-free SNP discovery and phylogenetics of hundreds of microbial genomes. PLoS One 2013, 8:e81760.
67. Hohl M, Ragan MA: Is multiple-sequence alignment required for accurate inference of phylogeny? Syst Biol 2007, 56:206-221.
68. Gnerre S, Maccallum I, Przybylski D, Ribeiro FJ, Burton JN, Walker BJ, Sharpe T, Hall G, Shea TP, Sykes S, Berlin AM, Aird D, Costello M, Daza R, Williams L, Nicol R, Gnirke A, Nusbaum C, Lander ES, Jaffe DB: High-quality draft assemblies of mammalian genomes from massively parallel sequence data. Proc Natl Acad Sci U S A 2011, 108:1513-1518.
69. Simpson JT, Wong K, Jackman SD, Schein JE, Jones SJ, Birol I: ABySS: a parallel assembler for short read sequence data. Genome Res 2009, 19:1117-1123.
70. Simpson JT, Durbin R: Efficient de novo assembly of large genomes using compressed data structures. Genome Res 2012, 22:549-556.
71. Miller JR, Delcher AL, Koren S, Venter E, Walenz BP, Brownley A, Johnson J, Li K, Mobarry C, Sutton G: Aggressive assembly of pyrosequencing reads with mates. Bioinformatics 2008, 24:2818-2824.
72. Pevzner PA, Tang H, Waterman MS: An Eulerian path approach to DNA fragment assembly. Proc Natl Acad Sci U S A 2001, 98:9748-9753.
73. Zerbino DR: Using the Velvet de novo assembler for short-read sequencing technologies. Curr Protoc Bioinformatics 2010, Chapter 11:Unit 1115.
74. Margulies M, Egholm M, Altman WE, Attiya S, Bader JS, Bemben LA, Berka J, Braverman MS, Chen YJ, Chen Z, Dewell SB, Du L, Fierro JM, Gomes XV, Godwin BC, He W, Helgesen S, Ho CH, Irzyk GP, Jando SC, Alenquer ML, Jarvie TP, Jirage KB, Kim JB, Knight JR, Lanza JR, Leamon JH, Lefkowitz SM, Lei M, Li J, et al: Genome sequencing in microfabricated high-density picolitre reactors. Nature 2005, 437:376-380.
75. Li R, Fan W, Tian G, Zhu H, He L, Cai J, Huang Q, Cai Q, Li B, Bai Y, Zhang Z, Zhang Y, Wang W, Li J, Wei F, Li H, Jian M, Li J, Zhang Z, Nielsen R, Li D, Gu W, Yang Z, Xuan Z, Ryder OA, Leung FC, Zhou Y, Cao J, Sun X, Fu Y, et al: The sequence and de novo assembly of the giant panda genome. Nature 2010, 463:311-317.
76. Kurtz S, Phillippy A, Delcher AL, Smoot M, Shumway M, Antonescu C, Salzberg SL: Versatile and open software for comparing large genomes. Genome Biol 2004, 5:R12.
77. Delcher AL, Kasif S, Fleischmann RD, Peterson J, White O, Salzberg SL: Alignment of whole genomes MUMMER. Nucl Acids Res 1999, 27:2369-2369.
78. Perna NT, Plunkett G 3rd, Burland V, Mau B, Glasner JD, Rose DJ, Mayhew GF, Evans PS, Gregor J, Kirkpatrick HA, Posfai G, Hackett J, Klink S, Boutin A, Shao Y, Miller L, Grotbeck EJ, Davis NW, Lim A, Dimalanta ET, Potamousis KD, Apodaca J, Anantharaman TS, Lin J, Yen G, Schwartz DC, Welch RA, Blattner FR: Genome sequence of enterohaemorrhagic Escherichia coli O157:H7. Nature 2001, 409:529-533.
79. Loman NJ, Misra RV, Dallman TJ, Constantinidou C, Gharbia SE, Wain J, Pallen MJ: Performance comparison of benchtop high-throughput sequencing platforms. Nat Biotechnol 2012, 30:434-439.
80. Tatusova T, Ciufo S, Fedorov B, O'Neill K, Tolstoy I: RefSeq microbial genomes database: new representation and annotation strategy. Nucleic Acids Res 2014, 42:D553-D559.
81. Bankevich A, Nurk S, Antipov D, Gurevich AA, Dvorkin M, Kulikov AS, Lesin VM, Nikolenko SI, Pham S, Prjibelski AD, Pyshkin AV, Sirotkin AV, Vyahhi N, Tesler G, Alekseyev MA, Pevzner PA: SPAdes: a new genome assembly algorithm and its applications to single-cell sequencing. J Comput Biol 2012, 19:455-477.
82. Koren S, Harhay GP, Smith TP, Bono JL, Harhay DM, McVey SD, Radune D, Bergman NH, Phillippy AM: Reducing assembly complexity of microbial genomes with single-molecule sequencing. Genome Biol 2013, 14:R101.
83. Delcher AL, Phillippy A, Carlton J, Salzberg SL: Fast algorithms for large-scale genome alignment and comparison. Nucleic Acids Res 2002, 30:2478-2483.
84. Deloger M, El Karoui M, Petit MA: A genomic distance based on MUM indicates discontinuity between most bacterial species and genera. J Bacteriol 2009, 191:91-99.
85. Edgar RC: MUSCLE: multiple sequence alignment with high accuracy and high throughput. Nucleic Acids Res 2004, 32:1792-1797.
86. Teacch D: Borrador sin revisión de traducción ni formato Contenidos. 1-341.
87. Bruen TC, Philippe H, Bryant D: A simple and robust statistical test for detecting the presence of recombination. Genetics 2006, 172:2665-2681.
88. Price MN, Dehal PS, Arkin AP: FastTree 2-approximately maximumlikelihood trees for large alignments. PLoS One 2010, 5:e9490.
89. Medini D, Donati C, Tettelin H, Masignani V, Rappuoli R: The microbial pan-genome. Curr Opin Genet Dev 2005, 15:589-594.
90. Harris SR, Feil EJ, Holden MT, Quail MA, Nickerson EK, Chantratita N, Gardete S, Tavares A, Day N, Lindsay JA, Edgeworth JD, de Lencastre H, Parkhill J, Peacock SJ, Bentley SD: Evolution of MRSA during hospital transmission and intercontinental spread. Science 2010, 327:469-474
91. van Vugt-Lussenburg BM, van der Weel L, Hagen WR, Hagedoorn PL: Biochemical similarities and differences between the catalytic [4Fe-4S] cluster containing fumarases FumA and FumB from Escherichia coli. PLoS One 2013, 8:e55549.
92. Eyre DW, Cule ML, Wilson DJ, Griffiths D, Vaughan A, O'Connor L, Ip CL, Golubchik T, Batty EM, Finney JM, Wyllie DH, Didelot X, Piazza P, Bowden R, Dingle KE, Harding RM, Crook DW, Wilcox MH, Peto TE, Walker AS: Diverse sources of C. difficile infection identified on whole-genome sequencing. N Engl J Med 2013, 369:1195-1205.
93. Koren S, Treangen TJ, Hill CM, Pop M, Phillippy AM: Automated ensemble assembly and validation of microbial genomes. BMC Bioinformatics 2014, 15:126.
94. Sebaihia M, Wren BW, Mullany P, Fairweather NF, Minton N, Stabler R, Thomson NR, Roberts AP, Cerdeno-Tarraga AM, Wang H, Holden MT, Wright A, Churcher C, Quail MA, Baker S, Bason N, Brooks K, Chillingworth T, Cronin A, Davis P, Dowd L, Fraser A, Feltwell T, Hance Z, Holroyd S, Jagels K, Moule S, Mungall K, Price C, Rabbinowitsch E, et al: The multidrug-resistant human pathogen Clostridium difficile has a highly mobile, mosaic genome. Nat Genet 2006, 38:779-786.
95. Dawson LF, Donahue EH, Cartman ST, Barton RH, Bundy J, McNerney R, Minton NP, Wren BW: The analysis of para-cresol production and tolerance in Clostridium difficile 027 and 012 strains. BMC Microbiol 2011, 11:86.
96. Young GP, Ward PB, Bayley N, Gordon D, Higgins G, Trapani JA, McDonald MI, Labrooy J, Hecker R: Antibiotic-associated colitis due to Clostridium difficile: double-blind comparison of vancomycin with bacitracin. Gastroenterology 1985, 89:1038-1045.
97. Venugopal AA, Johnson S: Current state of Clostridium difficile treatment options. Clin Infect Dis 2012, 55:S71-S76.
98. Comas I, Coscolla M, Luo T, Borrell S, Holt KE, Kato-Maeda M, Parkhill J, Malla B, Berg S, Thwaites G, Yeboah-Manu D, Bothamley G, Mei J, Wei L, Bentley S, Harris SR, Niemann S, Diel R, Aseffa A, Gao Q, Young D, Gagneux S: Out-of-Africa migration and Neolithic coexpansion of Mycobacterium tuberculosis with modern humans. Nat Genet 2013, 45:1176-1182.
99. Robinson DF, Foulds LR: Comparison of Phylogenetic Trees. Math Biosci 1981, 53:131-147.
100. Cock PJ, Fields CJ, Goto N, Heuer ML, Rice PM: The Sanger FASTQ file format for sequences with quality scores, and the Solexa/Illumina FASTQ variants. Nucleic Acids Res 2010, 38:1767-1771.
101. Danecek P, Auton A, Abecasis G, Albers CA, Banks E, DePristo MA, Handsaker RE, Lunter G, Marth GT, Sherry ST, McVean G, Durbin R: Genomes Project Analysis G: The variant call format and VCFtools. Bioinformatics 2011, 27:2156-2158.
102. Hsi-Yang Fritz M, Leinonen R, Cochrane G, Birney E: Efficient storage of high throughput DNA sequencing data using reference-based compression. Genome Res 2011, 21:734-740.
103. Lipman D, Flicek P, Salzberg S, Gerstein M, Knight R: Closure of the NCBI SRA and implications for the long-term future of genomics data storage. Genome Biol 2011, 12:402.
104. Korlach J, Bjornson KP, Chaudhuri BP, Cicero RL, Flusberg BA, Gray JJ, Holden D, Saxena R, Wegener J, Turner SW: Real-time DNA sequencing from single polymerase molecules. Methods Enzymol 2010, 472:431-455.
105. Matsen FA, Kodner RB, Armbrust EV: pplacer: linear time maximumlikelihood and Bayesian phylogenetic placement of sequences onto a fixed reference tree. BMC Bioinformatics 2010, 11:538.
106. Loytynoja A, Vilella AJ, Goldman N: Accurate extension of multiple sequence alignments using a phylogeny-aware graph algorithm. Bioinformatics 2012, 28:1684-1691.
107. Tenover FC, Arbeit RD, Goering RV, Mickelsen PA, Murray BE, Persing DH, Swaminathan B: Interpreting chromosomal DNA restriction patterns produced by pulsed-field gel electrophoresis: criteria for bacterial strain typing. J Clin Microbiol 1995, 33:2233-2239.
108. Maiden MC, Bygraves JA, Feil E, Morelli G, Russell JE, Urwin R, Zhang Q, Zhou J, Zurth K, Caugant DA, Feavers IM, Achtman M, Spratt BG: Multilocus sequence typing: a portable approach to the identification of clones within populations of pathogenic microorganisms. Proc Natl Acad Sci U S A 1998, 95:3140-3145.
109. Castillo-Ramirez S, Corander J, Marttinen P, Aldeljawi M, Hanage WP, Westh H, Boye K, Gulay Z, Bentley SD, Parkhill J, Holden MT, Feil EJ: Phylogeographic variation in recombination rates within a global clone of methicillin-resistant Staphylococcus aureus. Genome Biol 2012, 13:R126.
110. Croucher NJ, Finkelstein JA, Pelton SI, Mitchell PK, Lee GM, Parkhill J, Bentley SD, Hanage WP, Lipsitch M: Population genomics of post-vaccine changes in pneumococcal epidemiology. Nat Genet 2013, 45:656-663.
111. Croucher NJ, Harris SR, Fraser C, Quail MA, Burton J, van der Linden M, McGee L, von Gottberg A, Song JH, Ko KS, Pichon B, Baker S, Parry CM, Lambertsen LM, Shahinas D, Pillai DR, Mitchell TJ, Dougan G, Tomasz A, Klugman KP, Parkhill J, Hanage WP, Bentley SD: Rapid pneumococcal evolution in response to clinical interventions. Science 2011, 331:430-434.
112. Hayashi K, Morooka N, Yamamoto Y, Fujita K, Isono K, Choi S, Ohtsubo E, Baba T, Wanner BL, Mori H, Horiuchi T: Highly accurate genome sequences of Escherichia coli K-12 strains MG1655 and W3110. Mol Syst Biol 2006, 2:2006.0007.
113. Rambaut A, Grassly NC: Seq-Gen: an application for the Monte Carlo simulation of DNA sequence evolution along phylogenetic trees. Comput Appl Biosci 1997, 13:235-238.
114. Parsnp github url. [https://github.com/marbl/parsnp/tree/master/script]
115. Yutin N, Galperin MY: A genomic update on clostridial phylogeny: Gram-negative spore formers and other misplaced clostridia. Environ Microbiol 2013, 15:2631-2641.

## doi:10.1186/s13059-014-0524-x

Cite this article as: Treangen et al.: The Harvest suite for rapid coregenome alignment and visualization of thousands of intraspecific microbial genomes. Genome Biology 2014 15:524.


[^0]:    * Correspondence: treangent@nbacc.net
    ${ }^{\dagger}$ Equal contributors
    National Biodefense Analysis and Countermeasures Center, 110 Thomas Johnson Drive, Frederick, MD 21702, USA

