Phables: from fragmented assemblies to high-quality bacteriophage genomes

Vijini Mallawaarachchi[1], Michael J. Roach[1], Bhavya Papudeshi[1], Sarah K. Giles[1], Susanna R. Grigson[1], Przemyslaw Decewicz[2,1], George Bouras[3], Ryan D. Hesse[1], Laura K. Inglis[1], Abbey L. K. Hutton[1], Elizabeth A. Dinsdale[1] and Robert A. Edwards[1]

[1]Flinders Accelerator for Microbiome Exploration, College of Science and Engineering, Flinders University, Bedford Park, Adelaide, SA, 5042, Australia

[2]Department of Environmental Microbiology and Biotechnology, Institute of Microbiology, Faculty of Biology, University of Warsaw, Warsaw 02-096, Poland

[3]Adelaide Medical School, The University of Adelaide, North Tce, Adelaide, SA, 5000, Australia

Keywords: bacteriophages, metagenomics, genomics, assembly graphs

Abstract

Motivation

Microbial communities found within the human gut have a strong influence on human health. Intestinal bacteria and viruses influence gastrointestinal diseases such as inflammatory bowel disease. Viruses infecting bacteria, known as bacteriophages, play a key role in modulating bacterial communities within the human gut. However, the identification and characterisation of novel bacteriophages remain a challenge. Available tools use similarities between sequences, nucleotide composition, and the presence of viral genes/proteins. Most available tools consider individual contigs to determine whether they are of viral origin. As a result of the challenges in viral assembly, fragmentation of viral genomes can occur, leading to the need for new approaches in viral identification.

Results

We introduce Phables, a new computational method to resolve bacteriophage genomes from fragmented viral metagenomic assemblies. Phables identifies bacteriophage-like components in the assembly graph, models each component as a flow network, and uses graph algorithms and flow decomposition techniques to identify genomic paths. Experimental results of viral metagenomic samples obtained from different environments show that over 80% of the bacteriophage genomes resolved by Phables have high quality and are longer than the individual contigs identified by existing viral identification tools.

Availability and Implementation

Phables is available on GitHub at https://github.com/Vini2/phables (DOI: 10.5281/zenodo.7645166).

Contact

vijini.mallawaarachchi@flinders.edu.au

Introduction

Bacteriophages (hereafter *phages*) are viruses that infect bacteria, which influence microbial ecology and help modulate microbial communities (Edwards and Rohwer 2005; Rodriguez-Valera et al. 2009). Phages are considered the most abundant biological entity on earth, totalling an estimated 10^{31} particles (Comeau et al. 2008). Since their discovery by Frederick Twort in 1915 (Twort 1915), phages have been isolated from many diverse environments (Keen 2015). When sequencing technologies were first developed, phage genomes were the first to be sequenced due to their relatively small genome size (Sanger et al. 1977). With the advent of second-generation sequencing technologies, the first metagenomic samples to be sequenced were phages (Breitbart et al. 2002). The availability of advanced sequencing technologies has facilitated the investigation of the effect of phages on the functions of microbial communities, especially in the human body's niche areas. For example, phages residing in the human gut have a strong influence on human health (Łusiak-Szelachowska et al. 2017) and impact gastrointestinal diseases such as inflammatory bowel disease (IBD) (Norman et al. 2015). To date, our understanding of the diversity of phages is limited, as most have not been cultured due to the inherent difficulty of recovering phages from their natural environments. Although countless millions of phage species are thought to exist, only 25,152 complete phage genomes have been sequenced according to the INfrastructure for a PHAge REference Database (INPHARED) (Cook et al. 2021) (as of the April 2023 update).

Metagenomics has enabled the application of modern sequencing techniques for the culture-independent study of microbial communities (Hugenholtz, Goebel, and Pace 1998). Metagenomic sequencing provides a multitude of sequencing reads from the genetic material in environmental samples that are composed of a mixture of prokaryotic, eukaryotic, and viral species. Metagenomic analysis pipelines start by assembling sequencing reads from metagenomic samples into longer contiguous sequences that are used in downstream analyses. Most metagenome assemblers (D. Li et al. 2015; Nurk et al. 2017) use graphs as the primary data structure to capture overlaps between reads. The *assembly graph* consists of vertices that represent fragments and directed edges that represent connections or overlaps between these fragments. Non-branching paths in the assembly graph (paths where all vertices have an in-degree and out-degree of one, except for the first and last vertices) are referred to as *unitigs* (Kececioglu and Myers 1995). Unitigs are entirely consistent with the read data and belong to the final genome(s). Assemblers such as SPAdes (Bankevich et al. 2012) condense

unitigs into single vertices and resolve longer optimised paths through the branches into contiguous sequences known as *contigs*. This can cause contigs to break when the assembler cannot confidently resolve a path through a branching point, resulting in under-assembly (Bankevich et al. 2012; Tomescu and Medvedev 2016). For metagenome assemblies, where a mixture of organisms is present in the assembly graph, under-assembly results in fragmentation, making it challenging to recover complete genomes. However, previous studies have shown that contigs that are connected to each other are more likely to belong to the same genome (V. Mallawaarachchi, Wickramarachchi, and Lin 2020; V. G. Mallawaarachchi, Wickramarachchi, and Lin 2020, 2021). Hence, the assembly graph retains important connectivity and neighbourhood information within fragmented assemblies. This concept has been successfully applied to develop tools such as GraphMB (Lamurias et al. 2022), MetaCoAG (V. Mallawaarachchi and Lin 2022a, 2022b), and RepBin (Xue et al. 2022), where the assembly graphs are utilised in conjunction with taxonomy-independent metagenomic binning methods to recover high-quality metagenome-assembled genomes (hereafter *MAGs*) of bacterial genomes. Moreover, assembly graphs have been used for bacterial strain resolution in metagenomic data (Quince et al. 2021). However, limited studies have been done to resolve phage genomes in metagenomic data.

Computational tools have enabled large-scale studies to recover novel phages entirely from metagenomic sequencing data (Simmonds et al. 2017) and gain insights into interactions with their hosts (Nayfach, Páez-Espino, et al. 2021; M. J. Roach, McNair, et al. 2022; Hesse et al. 2022). While exciting progress has been made towards identifying new phages, viral dark matter remains vast. Current methods are either too slow or result in inaccurate or incomplete phage genomes. Generating high-quality phage genomes via *de novo* metagenome assembly is challenging due to the modular and mosaic nature of phage genomes (Lima-Mendez, Toussaint, and Leplae 2011; Hatfull 2008; Belcaid, Bergeron, and Poisson 2010). Repeat regions result in fragmented assemblies and chimeric contigs (Casjens and Gilcrease 2009; Merrill et al. 2016). Hence, current state-of-the-art computational tools rely on the combination of either more conservative tools based on sequence- and profile-based screening (e.g. MetaPhinder (Jurtz et al. 2016)) or machine learning approaches based on nucleotide signatures (e.g. Seeker (Auslander et al. 2020), refer to Table S1 in section 1 of the Supplementary material). Resulting predictions are then evaluated using tools such as CheckV (Nayfach, Camargo, et al. 2021) and VIBRANT (Kieft, Zhou, and Anantharaman 2020) to categorise the predicted phages based on their completeness, contamination levels, and possible lifestyle (virulent or temperate) (McNair, Bailey, and Edwards 2012). Due to the

supervised nature of the underlying approaches, most of these tools cannot characterise novel viruses that are significantly different from known viruses. Moreover, the approach used by these tools can be problematic with fragmented assemblies where contigs do not always represent complete genomes. In an attempt to address this limitation, tools such as MARVEL (Amgarten et al. 2018) and PHAMB (Johansen et al. 2022) were developed to identify viral metagenome-assembled genomes (vMAGs) of phages from metagenomic data. These programs are reliant on existing taxonomy-independent metagenomic binning tools such as MetaBAT2 (Kang et al. 2019) or VAMB (Nissen et al. 2021) and attempt to predict viral genome bins from this output using machine learning techniques.

Metagenomic binning tools are designed to capture nucleotide and sequence coverage-specific patterns of different taxonomic groups; therefore, sequences from viruses with low and uneven sequence coverage are often improperly binned. Many metagenomic binning tools filter out short sequences (e.g., shorter than 1,500 bp (Kang et al. 2019)), which further result in the loss of essential regions in phage genomes that are often present as short fragments in the assembly (Casjens and Gilcrease 2009). Moreover, most metagenomic binning tools struggle to distinguish communities of similar viral variants (viral quasispecies), and bins often contain a mixture of multiple strains resulting in poor-quality MAGs (Nissen et al. 2021). Existing solutions that are developed to perform viral quasispecies assembly only consider one species at a time (Baaijens, Stougie, and Schönhuth 2020; Freire et al. 2021, 2022) and cannot be applied to complex metagenomes. Despite the recent progress, it is challenging for currently available tools to recover complete high-quality phage genomes from metagenomic data, and a novel approach is required to address this issue. The use of connectivity information from assembly graphs could potentially help overcome these challenges (as shown in previous studies on bacterial metagenomes (V. Mallawaarachchi and Lin 2022a, 2022b; Lamurias et al. 2022)) to enable the recovery of high-quality phage genomes.

In this paper, we introduce Phables, a software tool that can resolve complete high-quality phage genomes from viral metagenome assemblies. Phables uses the assembly graph and read mapping information to resolve phage-like components into genomic paths representing trustworthy variations of closely-related phages using graph algorithms and flow decomposition techniques. We evaluate the quality of the resolved genomes using different assessment techniques and demonstrate that Phables produces complete and high-quality phage genomes.

Materials and Methods

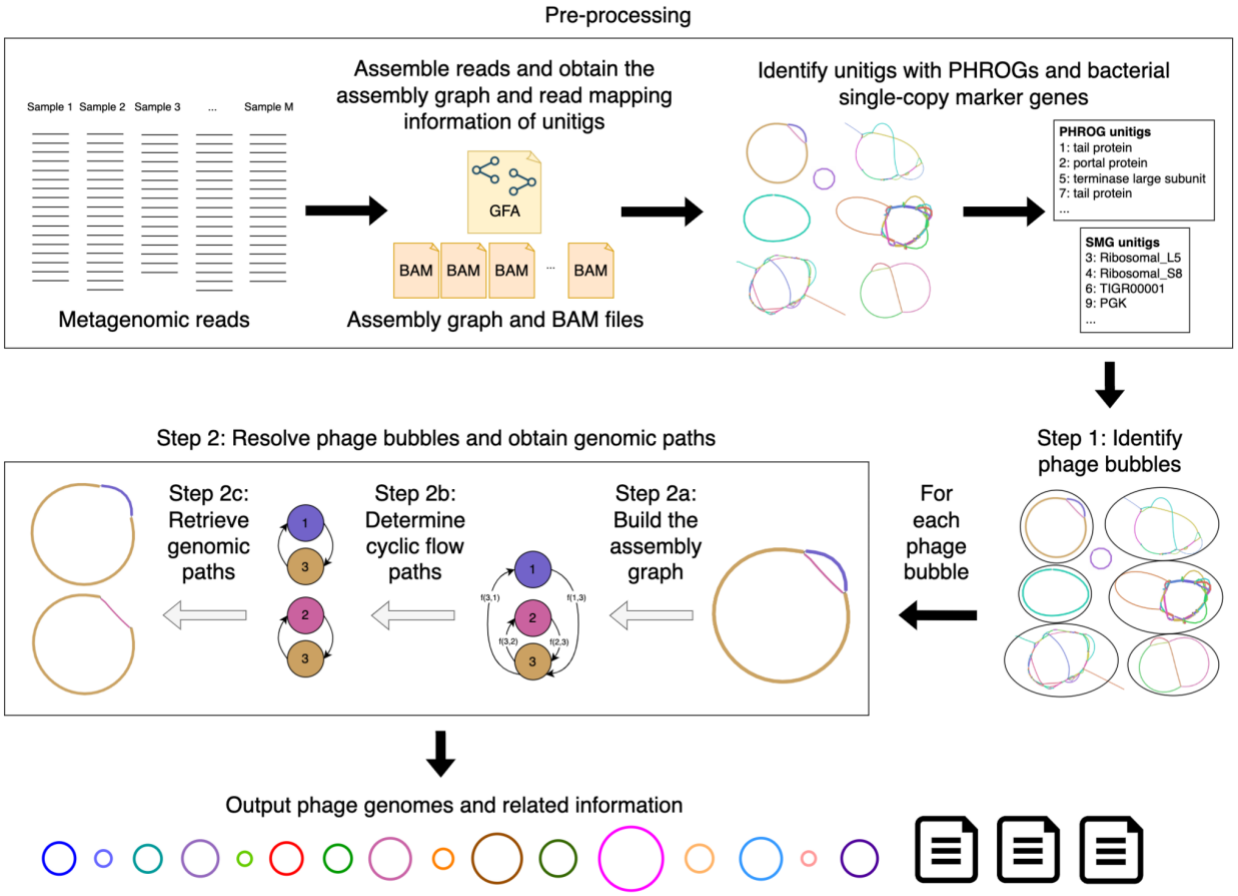

Figure 1: Phables workflow. Pre-processing: Assemble reads, obtain the assembly graph and read mapping information, and identify unitigs with PHROGs and bacterial single copy markers. Step 1: Identify and visualise phage bubbles found in the initial assembly. Step 2: For each phage bubble, build the assembly graph, determine the cyclic flow paths and retrieve the genomic paths. Finally, output phage genomes and related information.

Here we present the overall workflow of Phables (Figure 1). Metagenomic reads from multiple environmental and human virome samples are assembled, and the assembly graph and read mapping information are obtained. The unitig sequences from the assembly graph are extracted and screened for Prokaryotic Virus Remote Homologous Groups (PHROGs) (Terzian et al. 2021) and bacterial single-copy marker genes. Phables identifies cyclic sub-graphs (known as *phage bubbles*) and resolves separate phage genomes from each phage bubble. Finally,

Phables outputs the resolved phage genomes and related information. Each step of Phables is explained in detail in the following sections.

Pre-processing

The pre-processing step performed by Phables uses an assembly graph and generates the read mapping information and the gene annotations required for Step 1 in the workflow. We recommend Hecatomb (M. J. Roach, Beecroft, et al. 2022) to assemble the reads into contigs and obtain the assembly graph. However, Phables will work with any assembly graph in Graphical Fragment Assembly (GFA) format.

The unitig sequences are extracted from the assembly graph, and the raw sequencing reads are mapped to the unitigs using Minimap2 (H. Li 2018) and Samtools (H. Li et al. 2009). Phables then calculates the read coverage of unitigs, using the reads from all samples, and records the average read depth of each unitig, in terms of reads per kilobase sequence per million mapped reads (RPKM).

Phables identifies unitigs containing Prokaryotic Virus Remote Homologous Groups (PHROGs) (Terzian et al. 2021). PHROGs are viral protein families commonly used to annotate prokaryotic viral sequences. MMSeqs2 (Steinegger and Söding 2017) is used to identify PHROGs in unitigs using an identity cutoff of 30% and an e-value of less than 10^{-10}.

Phables identifies unitigs containing bacterial single-copy marker genes. Most bacterial genomes have conserved genes known as single-copy marker genes (SMGs) that appear only once in a genome (Dupont et al. 2012; Albertsen et al. 2013). FragGeneScan (Rho, Tang, and Ye 2010) and HMMER (Eddy 2011) are used to identify sequences containing SMGs. SMGs are considered to be present if more than 50% (by default) of the gene length is aligned to the unitig. The list of SMGs is provided in Table S2 in section 2 of the Supplementary material.

Step 1: Identify phage bubbles

Phables identifies cyclic components from the final assembly graph which contain unitigs that do not have any bacterial SMGs (identified from the preprocessing step) and contain one or more genes belonging to a PHROG for at least one of the PHROG categories: *head and packaging, connector, tail* and *lysis* which contain known phage markers and are highly conserved in tailed phages (Auslander et al. 2020). The presence of selected PHROGs ensures the components

are phage-like and represent potential phage genomes. The absence of bacterial single-copy marker genes further ensures that the components are not prophages or plasmids. However, it must be noted that Phables reports all unitigs in components containing phage-like proteins, which allows users for further evaluation, in particular the ones carrying prophages.

These identified components are referred to as *phage bubbles* due to their appearance when visualising the assembly graph. Components that are comprised of a single circular unitig (the two ends of the unitig overlap) and that satisfy the above conditions for genes are considered *phage bubbles* only if the circular unitig is longer than the predefined threshold $minlength$ that is set to 2,000 bp by default, as this is approximately the lower bound of genome length for tailed phages (Luque et al. 2020).

Step 2: Build assembly graphs of phage bubbles and obtain genomic paths

Step 2a: Build the assembly graph

Following the definitions from STRONG (Quince et al. 2021), we define the assembly graph $G = (V, E)$ for a phage bubble where $V = \{1, 2, 3, , |V|\}$ is a collection of unitig sequence vertices that make up a phage bubble and directed edges $E \in V \times V$ represent connections between unitigs. Each directed edge $(u^{d_1} \rightarrow v^{d_2})$ is defined by a starting vertex u and an ending vertex v (the arrow denotes the direction of the overlap), where $d_1, d_2 \in \{+, -\}$ indicates whether the overlap occurs between the original sequence, indicated by a $+$ sign or its reverse complement, indicated by a – sign.

The *weight* of each edge $(u^{d_1} \rightarrow v^{d_2})$ irrespective of the orientation of the edge, termed $w(u \rightarrow v)$ is set to the minimum of the read coverage values of the two unitigs u and v. We also define the *confidence* of each edge $(u^{d_1} \rightarrow v^{d_2})$ irrespective of whether the overlap occurs between the original sequence or its reverse complement, termed $conf(u \rightarrow v)$ is defined as the number of paired-end reads spanning across $(u \rightarrow v)$. Here, the forward read maps to unitig u and the reverse read maps to unitig v. Paired-end information has been used in previous studies for assembling viral quasispecies (Freire et al. 2022) to untangle assembly graphs. The more paired-end reads map to the pair of unitigs, the more confident we are about the overlap represented by the edge.

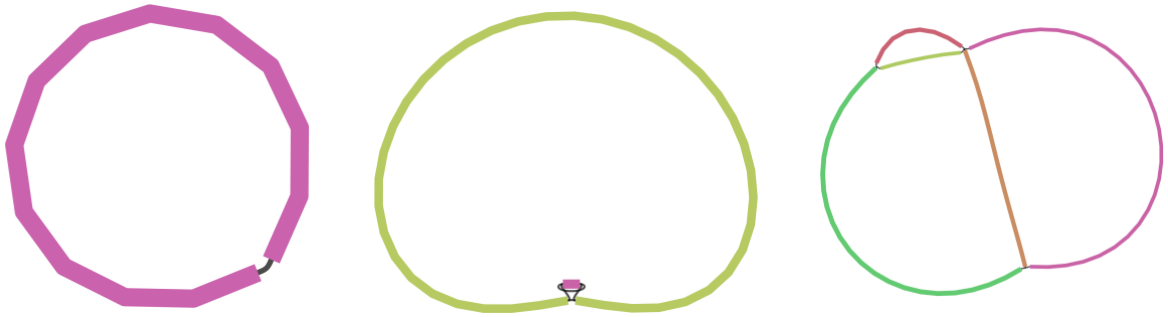

(a) Case 1 (b) Case 2 (c) Case 3

Figure 2: Cases of phage bubbles. (a) Case 1 represents a phage bubble with one circular unitig and is considered a complete phage genome. (b) Case 2 represents a phage bubble with two circular unitigs connected to each other. (c) Case 3 represents a phage bubble that is more complex with multiple unitigs and multiple paths.

Step 2b: Determine cyclic flow paths

We model the graph of the phage bubble as a flow decomposition problem and obtain the cyclic paths with their coverage values calculated from the read coverages of unitigs and read mapping information. We define three cases based on the number of unitigs present in the phage bubble as shown in Figure 2. Each case will be discussed in detail in the following sub-sections.

Case 1: Phage bubble consists of one circular unitig

When the phage bubble has only one circular unitig longer than the predefined threshold $minlength$, Phables considers this unitig as one genome. The genomic path is defined as the unitig sequence itself.

Case 2: Phage bubble consists of two circular unitigs

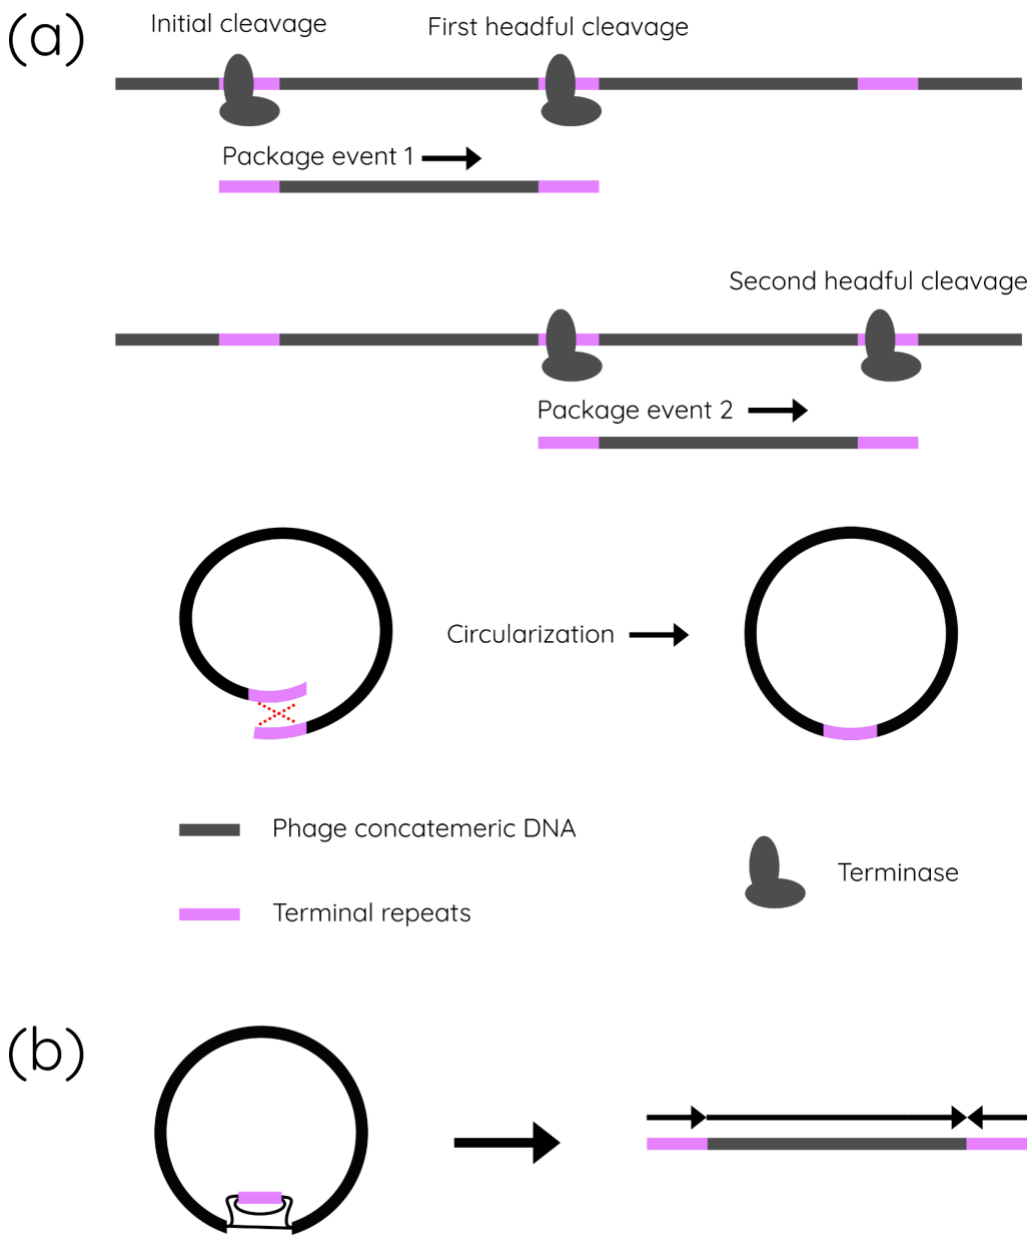

Figure 3: Example of resolving a case 2 phage bubble. (a) Each packaging event starts and finishes at a packaging recognition site that gets duplicated. The duplicated regions recombine and collapse during circularisation to form a continuous circular genome. (b) Representation of a case 2 phage bubble with a terminal repeat assembled separately and then resolved into a linearised genome.

The phage bubbles in case 2 have two circular unitigs connected together where at least one is longer than the predefined threshold $minlength$ and the shorter unitig is a few hundred base pairs long (generally shorter than 1,000 bp). This is an interesting case as this type of bubble structure describes the terminal repeats of phages. Some phages have double-stranded repeats at their termini which are a few hundred base pairs in length and are exactly the same in every virion chromosome (i.e. they are not permuted) (Casjens and Gilcrease 2009). The terminal repeats are generated by a duplication of the repeat region in concert with packaging (Zhang and Studier 2004; Yeon-Bo Chung et al. 1990) (Figure 3 (a)). This type of end structure could be overlooked when a phage genome sequence is determined by shotgun methods because sequence assembly can merge the two ends to give a circular sequence. Phables successfully resolves the terminal repeats to form complete genomes.

To resolve the phage bubble in case 2, we consider the shorter unitig (shorter than $minlength$) as the terminal repeat. Now we combine the original sequence of the terminal repeat to the beginning of the longer unitig and the reverse complement of the terminal repeat to the end of the longer unitig (Figure 3 (b)). The coverage of the path will be set to the coverage of the longer unitig.

Case 3: Phage bubble consists of three or more unitigs

In case 3, we have more complex phage bubbles where there are more than two unitigs forming branching paths, and we model them as a minimum flow decomposition (MFD) problem. The MFD problem decomposes a directed acyclic graph (DAG) into a minimum number of source-to-sink ($s-t$) paths that explain the flow values of the edges of the graph (Vatinlen et al. 2008; Dias et al. 2022). The most prominent applications of the MFD problem in bioinformatics include reconstructing RNA transcripts (Shao and Kingsford 2017; Tomescu et al. 2013; Gatter and Stadler 2019) and viral quasispecies assembly (Baaijens, Stougie, and Schönhuth 2020). The MFD problem is solved using integer linear programming (ILP) (Schrijver 1998). Since we are dealing with a similar problem to viral quasispecies assembly but specifically for phage genomes in metagenomes, Phables models each of the remaining phage bubbles as an MFD problem and uses the MFD-ILP implementation from Dias *et al.* (Dias et al. 2022). MFD-ILP finds a $FD(\mathcal{P}, w)$ with a set of $s-t$ flow paths $\mathcal{P}$ and associated weights w such that the number of flow paths is minimized. These flow paths represent possible genomic paths. An example of a phage bubble with possible paths is shown in Figure 4.

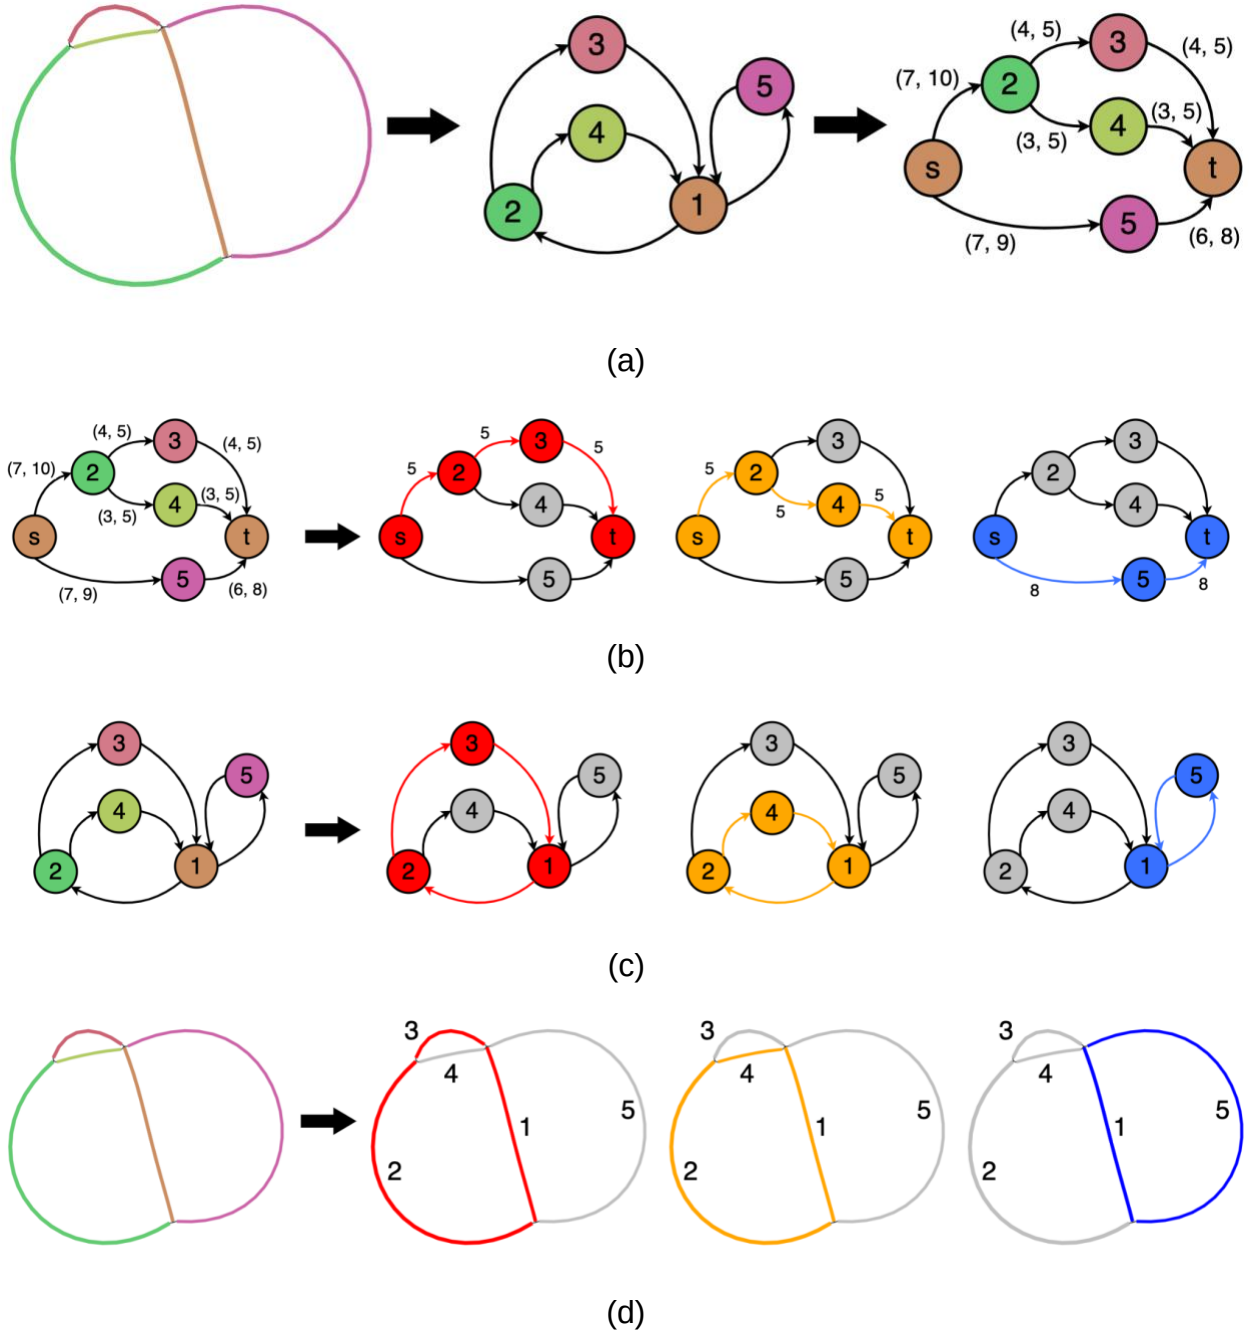

Figure 4: Example of a phage bubble (a) being modelled as a flow network and resolved into paths denoted using (b) flow network visualisation with flow values, (c) graph visualisation with directed edges and (d) Bandage (Wick et al. 2015) visualisation (with corresponding unitig numbers). Here three s-t flow paths (1→2→3, 1→2→4 and 1→5) can be obtained corresponding to three phage genomes. The arrows in (b)-(c) denote the resolution into 3 paths.

First, we convert the assembly graph of the phage bubble into a DAG. We start by removing dead-ends from G. We consider a vertex to be a dead-end if it has either no incoming edges or no outgoing edges, which arise due to errors at the start or end of reads that can create protruding chains of deviated edges (Bankevich et al. 2012). Dead-ends are particularly problematic in later steps of Phables as they can affect the continuity of circular genomic paths. Hence, their removal ensures that all the possible paths in the graph form closed cycles. We eliminate dead-ends by recursively removing vertices with either no incoming edges or no outgoing edges. Note that removing one dead-end can cause another vertex that is linked only to the removed one to become a dead-end, hence the removal process is done recursively.

Since a phage bubble forms a cyclic graph as shown in Figure 4 (a), we have to identify a vertex to represent the source/sink (referred to as st) in order to convert the graph to a DAG and model it as a flow network. Starting from every vertex ($source$), we conduct a breadth-first-search and identify an iterator, $(level, vertices)$, where $vertices$ is the non-empty list of vertices at the same distance $level$ from the $source$. The method that generates this iterator is known as bfs_layers and we use the NetworkX implementation (https://networkx.org/documentation/stable/reference/algorithms/generated/networkx.algorithms.traversal.breadth_first_search.bfs_layers.html#networkx.algorithms.traversal.breadth_first_search.bfs_layers). We extract the vertices in the final layer and check if their successors are equal to $source$. If this condition holds for some vertex in G, we consider this vertex to be the st vertex of G. If more than one vertex satisfies the condition to be a st vertex, then we pick the vertex corresponding to the longest unitig as the st vertex. This process is carried out to find a vertex common to the flow paths (refer to Algorithm S1 in section 3 of the Supplementary material). As an example, consider vertex 1 in Figure 4 (a). When we do a breadth-first-search starting from vertex 1, the vertices in the last layer in our iterator will be 3 and 4. The successor of both 3 and 4 is vertex 1. Since the successors of the vertices in the last layer are the same as the starting vertex, we consider vertex 1 as the st vertex.

The edges of G that are weighted according to unitig coverage, may not always satisfy the conservation of flow property because of uneven sequencing depths at different regions of the genomes (Peng et al. 2012; Gunasekera et al. 2021). Hence, we use inexact flow networks which allow the edge weights to belong to an interval. Once we have identified a st vertex, we separate that vertex into two separate verticies for the source s and sink t. We create an inexact flow network $G_f = (V, E, f, \bar{f})$ from s to t and model the rest of the vertices and edges in G. For example, in Figure 4 (b) vertex 1 is broken into two vertices s and t, and the network flows from s

to t,. For every edge $(u, v) \in E$ we have associated two positive integer values $f_{uv} \in f$ and $\bar{f_{uv}} \in \bar{f}$, satisfying $f_{uv} \leq \bar{f_{uv}}$ where $f_{uv} = conf(u \rightarrow v)$ and $\bar{f_{uv}} = w(u \rightarrow v)$. In Figure 4 (b), each edge has two values $(f_{uv}, \bar{f_{uv}})$ that define the flow interval for the inexact flow network G_f. By setting f_{uv} to $conf(u \rightarrow v)$ we ensure that the flow through the edge is bounded by the number of paired-end reads spanning across the edge. For example, in Figure 4 (b), the edge $(2 \rightarrow 3)$ has 4 paired-end reads spanning across the corresponding unitigs of 2 and 3 and a weight of 5 (which is the minimum of the read coverage values of the two unitigs 2 and 3 obtained from Step 2a). Hence, we set $f_{uv} = 4$ and $\bar{f_{u,v}} = 5$.

Next, we define a set of simple paths $\mathcal{R} = \{R_1, R_2, R_3, ..., R_l\}$ where the edges that form each path have paired-end reads spanning across them. Enforcing these paths to contain sequencing reads ensures that genuine connections are identified and reflected in at least one decomposed path. For example, in Figure 4 (b), the edge $(2 \rightarrow 3)$ has 4 paired-end reads spanning across the edge. Hence, we add the path $R_1 = (2, 3)$ to $\mathcal{R}$.

Now we model our inexact flow network G_f as a minimum inexact flow decomposition (MIFD) problem and determine a minimum-sized set of $s - t$ paths $\mathcal{P} = (P_1, P_2, P_3, ..., P_k)$ and associated weights $w = (w_1, w_2, w_3, ..., w_k)$ with each $w_i \in \mathbb{Z}^+$ where the following conditions hold.

1. $f_{uv} \leq \sum_{i \in \{1,...,k\} s.t. (u,v) \in P_i} w_i \leq \bar{f_{uv}} \quad \forall (u, v) \in E$
2. $\forall R_j \in \mathcal{R}, \exists P_i \in \mathcal{P} \text{ such that } R_j \text{ is a subpath of } P_i$

A path P_i will consist of unitigs with orientation information. The weight w_i will be the coverage of the genome represented by the path P_i.

Step 2c: Retrieve genomic paths

The flow paths obtained from cases 1 and 2 described in the previous section are directly translated to genomic paths based on the unitig sequences. In case 3, we get $s - t$ paths from the flow decomposition step (as shown in Figure 4 (b)). The paths longer than the predefined threshold $minlength$ and have a coverage value of $mincov$ (10 by default) or above are retained. For each remaining path, we remove t from the path as s and t are the same vertex and combine the nucleotide sequences of the unitigs corresponding to the vertices and the orientation of edges in the flow path (refer to Figure 4 (c) and (d)). Once the genomic paths of

phage bubbles are obtained, we record the constituent unitigs, path length (in bp), coverage (i.e., the flow value of the path) and the GC content of each genomic path.

Experimental Design

Datasets

We tested Phables on the following real viral metagenomic datasets available from the National Center for Biotechnology Information (NCBI).

1. Water samples from Nansi Lake and Dongping Lake in n Shandong Province, China (NCBI BioProject number PRJNA756429), referred to as **Lake Water**
2. Soil samples from flooded paddy fields from Hunan Province, China (NCBI BioProject number PRJNA866269), referred to as **Paddy Soil**
3. Wastewater virome (NCBI BioProject number PRJNA434744), referred to as **Wastewater**
4. Stool samples from patients with IBD and their healthy household controls (NCBI BioProject number PRJEB7772) (Norman et al. 2015), referred to as **IBD**

All the datasets were processed using Hecatomb version 1.0.1 to obtain a single assembly graph for each dataset (M. J. Roach, Beecroft, et al. 2022). Tables S3 and S4 in section 4 of the Supplementary material summarise the information about the datasets and their assemblies.

Tools benchmarked

We benchmarked Phables with PHAMB (Johansen et al. 2022), a recently published viral identification tool that predicts whether MAGs represent phages and outputs genome sequences. PHAMB takes binning results from a metagenomic binning tool and predicts bins that contain bacteriophage sequences. The MAGs for PHAMB were obtained by running VAMB (version 3.0.8), a binning tool that does not rely on bacterial marker genes, in co-assembly mode on the original contigs with author-recommended parameter `--minfasta 2000` and the `--cuda` flag. The commands used to run all the tools are found in section 5 of the Supplementary material.

Evaluation criteria

Evaluation criteria for binning tools

The resolved genomes from Phables and identified MAGs from PHAMB were evaluated using CheckV version 1.0.1 (Nayfach, Camargo, et al. 2021) (with reference database version 1.5) which compares bins against a large database of complete viral genomes. We compare the following metrics from the CheckV results.

1. CheckV viral quality
2. Completeness of sequences - number of sequences with >90% completeness
3. Contamination of sequences - number of sequences with <10% contamination
4. The number and length distribution of sequences with the following warnings
    a. Contig >1.5x longer than expected genome length
    b. High kmer_freq may indicate a large duplication

Since PHAMB predicts all viral bins, we only consider the bins from PHAMB that contain the contigs corresponding to the unitigs recovered by Phables for a fair comparison.

Evaluation criteria for resolved genomes

The number of components resolved by Phables for each case were recorded. The viral quality of the resolved genomes and the unitigs and contigs contained in the corresponding genomic paths were evaluated using CheckV version 1.0.1 (Nayfach, Camargo, et al. 2021).

Results

Benchmarking results

Phables resolves unitigs within phage bubbles to produce multiple complete and high-quality genomes from the viral metagenomes (Figure 5). The genome quality of Phables results was compared with the viral-MAG prediction tool PHAMB (Johansen et al. 2022) and evaluated using CheckV (Nayfach, Camargo, et al. 2021). Figure 5 denotes the comparison of length distributions and genome/bin counts of different CheckV quality categories for Phables and PHAMB results. Unlike Phables, PHAMB has produced genomes with longer sequences shown in Figures 5 (a), (c), (e) and (g), because PHAMB combines all the contigs in a bin to form one

long sequence. As denoted in Figures 5 (b), (d), (f) and (h), Phables has recovered the most number of complete and high-quality genomes combined for all the datasets; 136 (86.08%) in Lake Water, 368 (90.42%) in Paddy Soil, 57 (64.04%) in Wastewater and 142 (91.61%) in IBD.

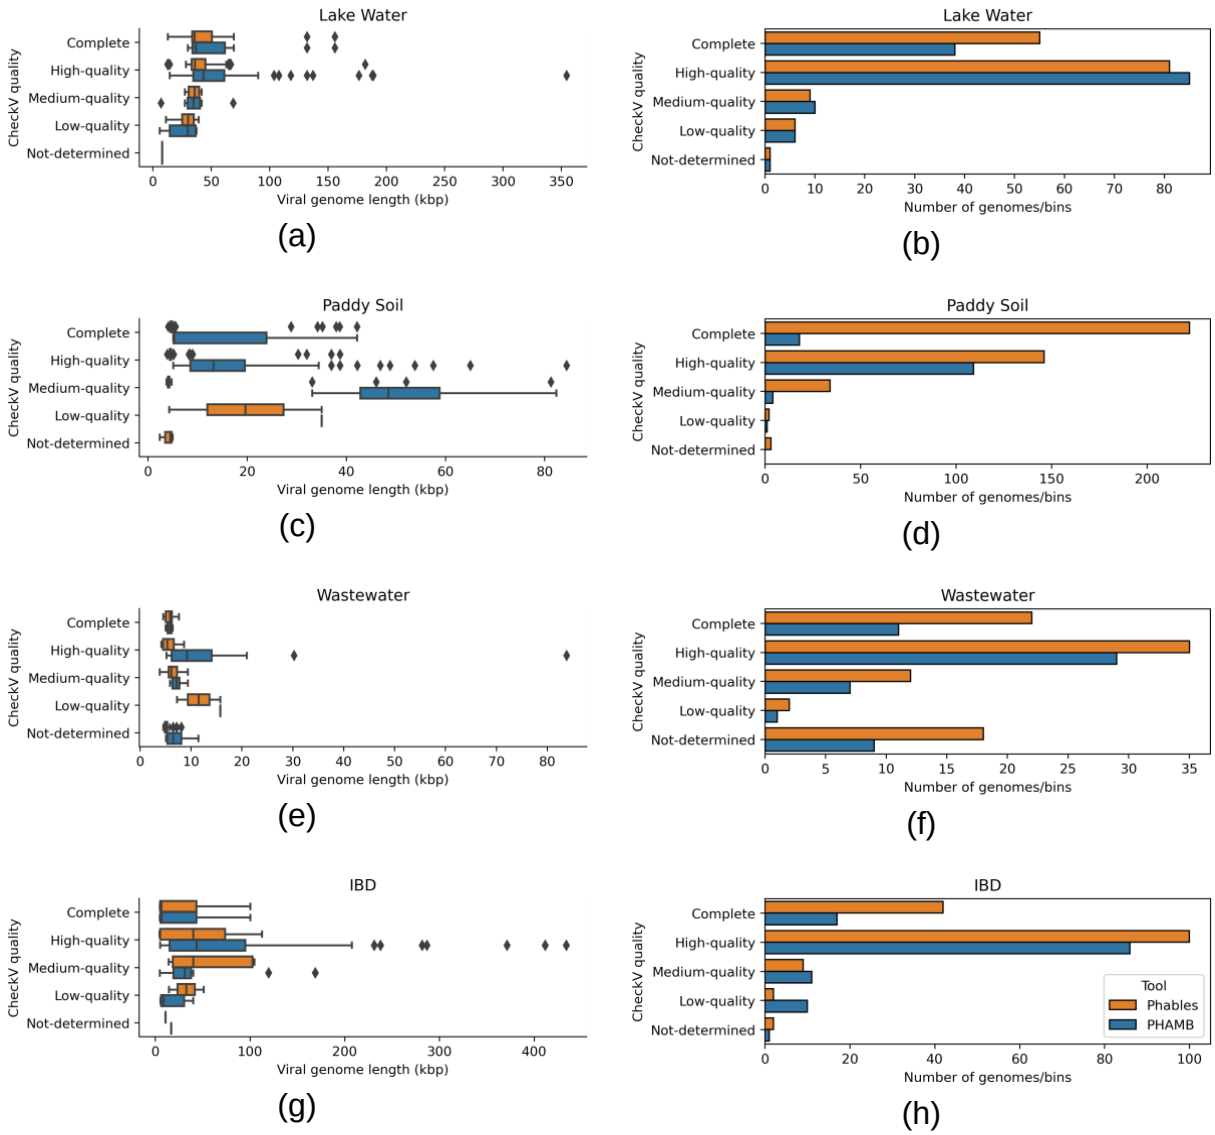

Figure 5: Length distribution and count of sequences from Phables and PHAMB (Johansen et al. 2022) having different CheckV quality categories for the viral metagenomic datasets (a) - (b) Lake Water, (c) - (d) Paddy soil, (e) - (f) Wastewater, and (g) - (h) IBD.

Phables accurately recovers short sequences (such as terminal repeats that are challenging for metagenomic binning tools to recover) using the assembly graph and produces high-quality genomes. We observed that VAMB miss-binned the majority of the short sequences, which reduced the quality of PHAMB results. For example, the repeat sequences in the case 2 bubbles identified by Phables had a mean length of 600 bp in Lake Water, 649 bp in Paddy Soil, 511 bp in Wastewater and 638 bp in IBD datasets (refer to Table S5 in section 6 of the Supplementary material for exact lengths of the sequences). All of these short sequences, except for those in the IBD dataset were found in a different bin than the bin of their connected longer sequence in the PHAMB results (8 out of 8 in Lake Water, 2 out of 2 in Paddy Soil and 1 out of 1 in Wastewater). Phables recovered these short repeat sequences along with their connected longer sequences within a phage bubble using the connectivity information of the assembly graph.

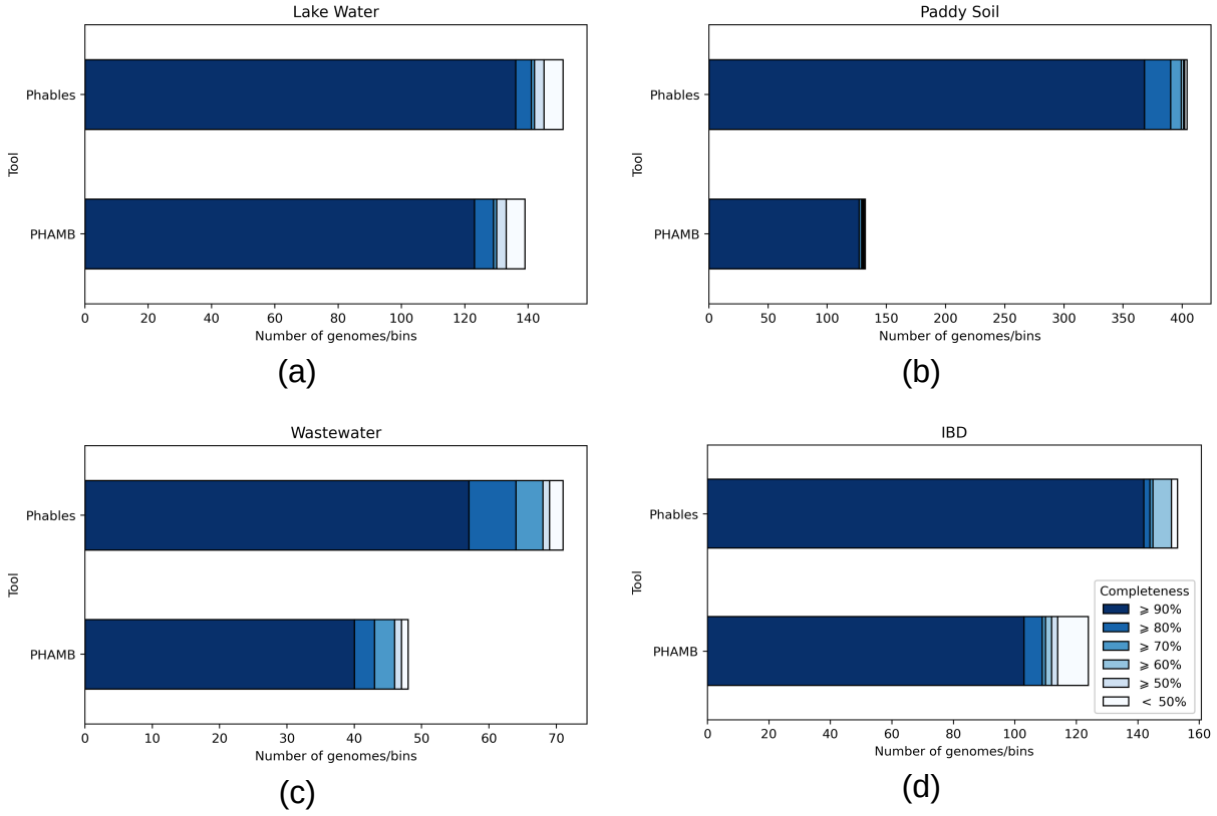

Figure 6: Comparison of CheckV completeness of sequences from Phables and PHAMB (Johansen et al. 2022) for the viral metagenomic datasets (a) Lake Water, (b) Paddy soil, (c) Wastewater, and (d) IBD.

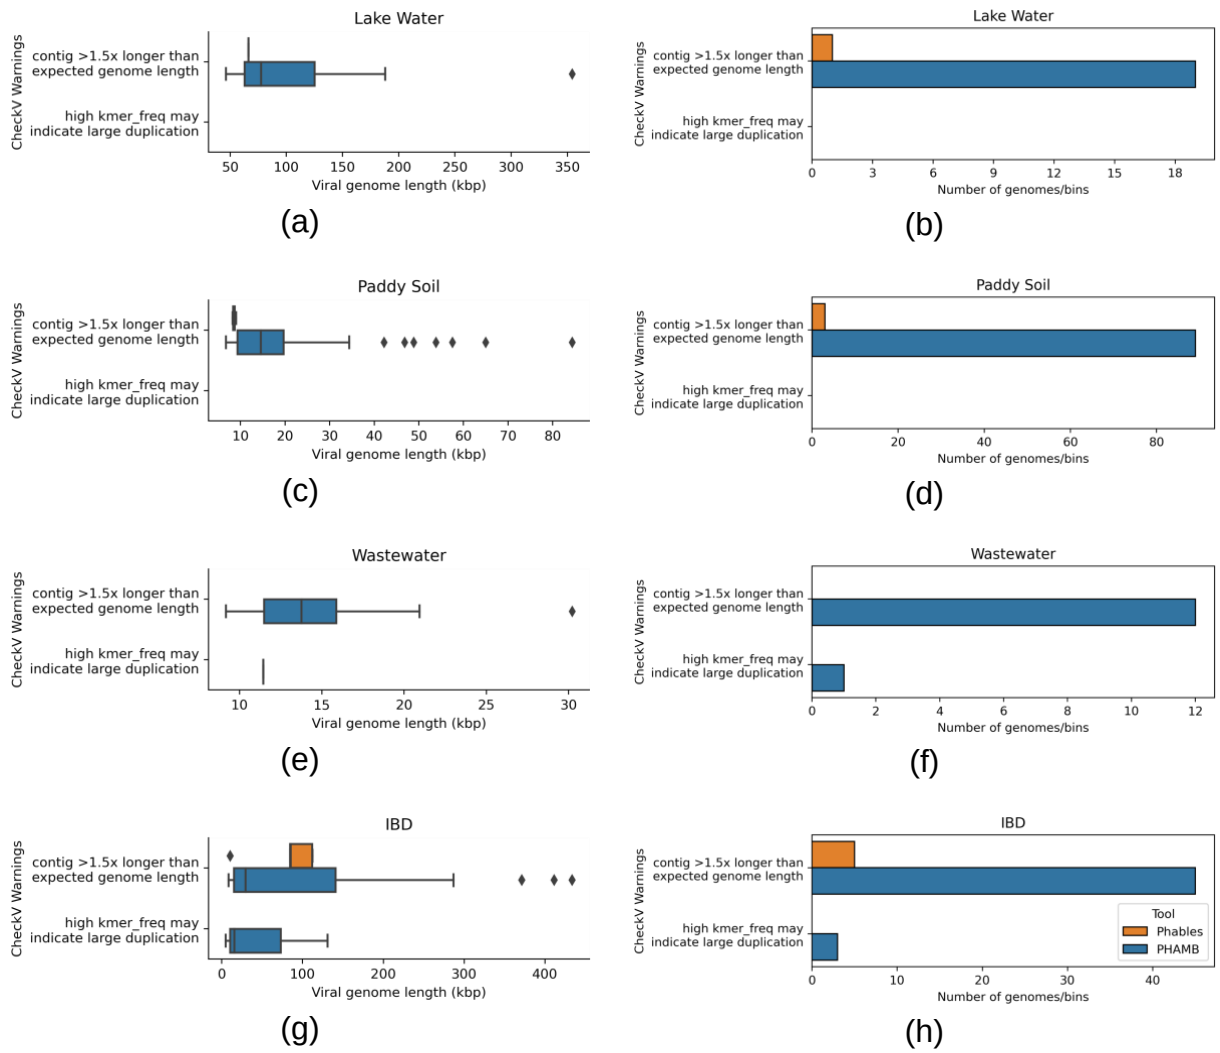

Figure 7: Length distribution and count of sequences having the selected CheckV warnings from Phables and PHAMB (Johansen et al. 2022) results of the viral metagenomic datasets (a) - (b) Lake Water, (c) - (d) Paddy soil, (e) - (f) Wastewater, and (g) - (h) IBD.

Phables improved the completeness results of genomes from all environments. Figure 6 demonstrates the number of genomes/bins with different ranges of CheckV completeness using the results from Phables and PHAMB. Phables has recovered the most genomes with over 90% completeness in all the datasets (136 for Lake Water, 368 for Paddy Soil, 57 for Wastewater and 142 for IBD), which consist of complete and high-quality genomes. Phables has been able

to recover genomes with high completeness because of the resolution steps that combine the fragmented sequences correctly.

Phables resulted in a high number of low-quality genomes as determined by CheckV in the Wastewater dataset compared to the other datasets (Figure 5 (f)). A possible reason for this is that these may be novel phages (as they contain conserved phage markers even though CheckV categorises them as “low-quality” or “not-determined”), and so they are not yet present in the databases that CheckV relies on.

PHAMB does not carry out any resolution steps when combining the contigs of identified MAGs, which results in high levels of erroneous genome structures, contamination and duplications within genomes because of the presence of multiple closely-related phages. Such duplications are identified from the warnings reported by CheckV. Hence, we evaluated the number and length distribution of sequences having CheckV warnings and the results are shown in Figure 7. PHAMB has produced the highest number of genomes with CheckV warnings and produced some very long genomes (~355 - 485 kbp as shown in Figures 7 (a) and (g)), suggesting the combination of two or more variants together in a bin. Only a few genomes produced from Phables (5 or less) contain CheckV warnings (refer to Table S6 in section 7 of the Supplementary material for the exact number of genomes with warnings for each dataset). These results show that Phables accurately recovers closely-related phages with mosaic genome structures (including regions like terminal repeats) from viral metagenomic samples and produces more high-quality/complete genomes compared to existing state-of-the-art tools.

Components resolved and quality comparison

The number of phage bubbles resolved by Phables under each case was recorded for all the datasets (Table 1). The majority of the resolved bubbles belong to either case 1 with a single circular unitig or case 2 with the terminal repeat. When resolving case 2 bubbles, Phables provides information regarding terminal repeats such as the length of the repeat region, that will be overlooked by other tools. Except for the IBD and Paddy Soil datasets, Phables was able to resolve all the case 3 phage bubbles from the rest of the datasets. In a few cases, the case 3 phage bubbles could not be resolved because Phables was unable to find a st vertex for these very complex bubbles (Refer to Figure S1 in section 8 of the Supplementary material for examples of unresolved phage bubbles).

Table 1. The number of components found and resolved for different cases of phage bubbles

Dataset	Case 1		Case 2		Case 3	
	Resolved	Found	Resolved	Found	Resolved	Found
Lake Water	135	135	8	8	6	6
Paddy Soil	401	401	2	2	3	5
Wastewater	85	85	1	1	1	1
IBD	99	99	2	2	22	32

Assemblers attempt to resolve longer paths in the assembly graph by connecting unitigs to form contigs (Bankevich et al. 2012; Kolmogorov et al. 2019). However, they are still unable to resolve complete genomes for complex datasets due to the mosaic nature of phage genomes and produce fragmented assemblies. Phables can be used to resolve these problematic contigs (or unitigs) and obtain high-quality genomes. Figure 8 denotes the comparison of CheckV quality of the genomes resolved in Phables and the unitigs and contigs included in the cases 2 and 3 phage bubbles. The most complete and high-quality sequences can be found as genomes (55 and 81 for Lake Water, 222 and 146 for Paddy Soil, 22 and 35 for Wastewater, and 42 and 100 for IBD, respectively). In contrast, most medium- and low-quality genomes can be found from contigs and unitigs. On average, more than 83% of the genomes resolved by Phables (from all cases) are either complete or have high quality. Hence, genomes resolved using Phables have higher quality and will be better candidates for downstream analysis than contigs.

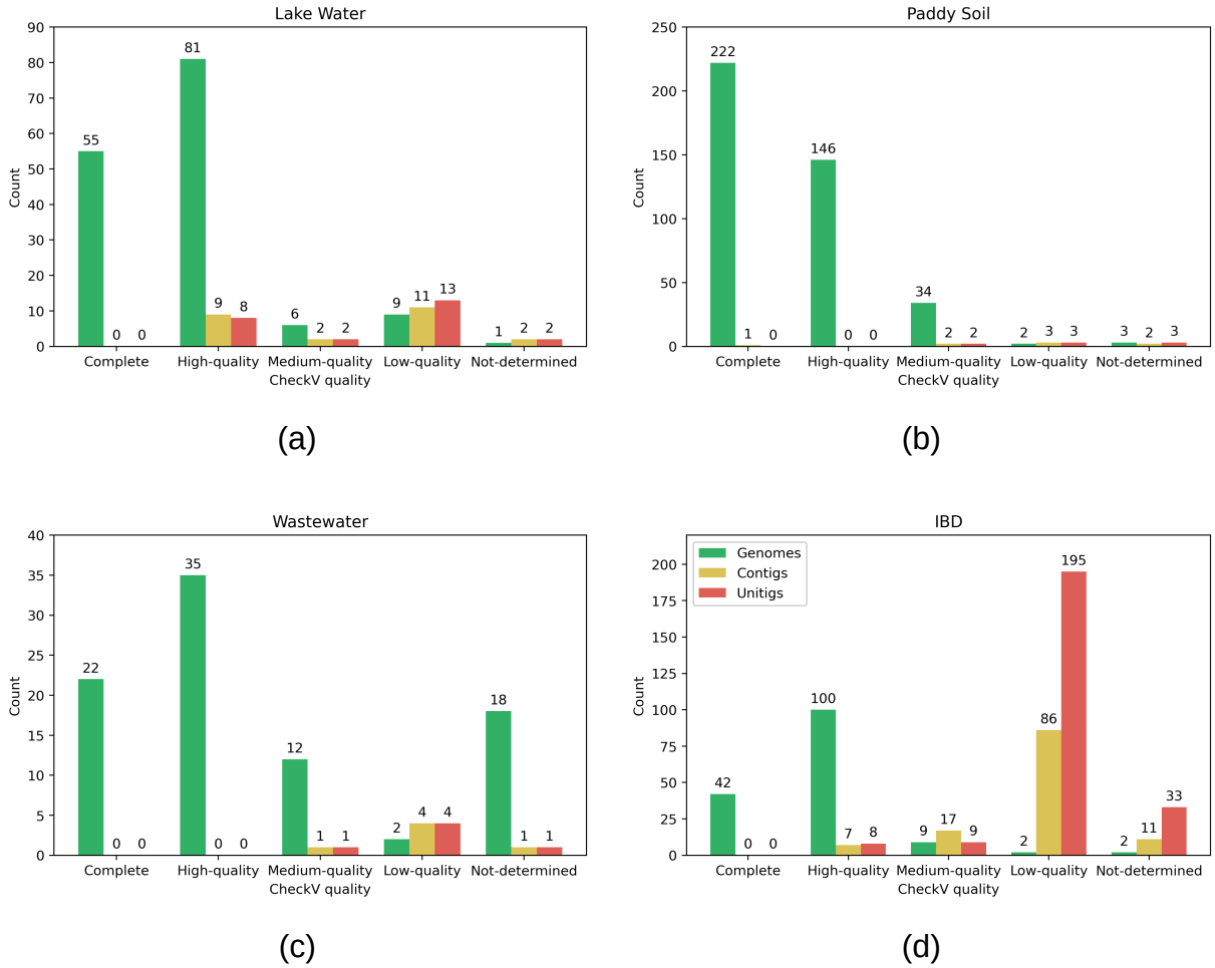

Figure 8: Counts of resolved genomes of Phables, unitigs and contigs included in case 2 and 3 phage bubbles with different CheckV qualities in the viral metagenomic datasets (a) Lake Water, (b) Paddy soil, (c) Wastewater, and (d) IBD.

Phage bubbles from other assembly methods

We further tested Phables with co-assemblies obtained from other metagenomic assemblers including metaSPAdes (Nurk et al. 2017) and MEGAHIT (D. Li et al. 2015) to show that the bubble-like components observed in the assembly graph are not an artefact of the assembly approach used in Hecatomb. Co-assembly is done by combining reads from multiple metagenomes and assembling them together, which increases the sequencing depth and provides sufficient coverage for low-abundance genomes to be recovered (Delgado and Andersson 2022). However, this becomes a computationally intensive approach as the number

of samples increases, and hence we have limited the results to just the Lake Water dataset. The results are provided in section 9 of the Supplementary material and show that the phage bubble structures are still present in the assemblies and were correctly resolved by Phables.

Implementation and resource usage

The source code of Phables was implemented using Python 3.10.8 and is available as a pipeline (including all the preprocessing steps) developed using Snaketool (M. Roach et al. 2022). The commands used to run all the software can be found in section 5 of the Supplementary material. The running times of Phables core methods and including the preprocessing steps were recorded for all the datasets and can be found in Tables S8 and S9 in section 10 of the Supplementary material. The core methods of Phables can be run in under 2 minutes with less than 4 gigabytes of memory for all the datasets.

Phables uses a modified version of the MFD-ILP implementation from Dias *et al.* (Dias et al. 2022) which supports inexact flow decomposition with subpath constraints. Gurobi version 10.0.1 was used as the ILP solver. To reduce the complexity of the ILP solver, the maximum number of unitigs in a phage bubble to be solved was limited to 200.

Discussion

The majority of the existing viral identification tools rely on sequence similarity- and profile-based approaches, only identifying whether assembled sequences are of viral origin, and cannot produce complete and high-quality phage genomes. Viral binning tools have been able to overcome these shortcomings up to a certain extent by producing viral MAGs, but these MAGs are fragmented and do not represent continuous genomes. Generally, the assembly process produces a large number of short contigs where some represent regions which while important are challenging to resolve in phages, such as terminal repeat regions. These short contigs are discarded or binned incorrectly by viral binning tools, producing incomplete MAGs. Moreover, the mosaic genome structure of phages results in viral communities that have closely related phages and phages that infect a common host, and assembling them can result in bubble structures which will be overlooked by existing binning tools. The resulting MAGs may contain multiple variants together and hence have high contamination.

Here, we introduce Phables, a new tool to resolve complete and high-quality phage genomes from viral metagenome assemblies using assembly graphs and flow decomposition techniques. We studied the assembly graphs constructed from different assembly approaches and different assembly software and consistently observed phage-like components with variation (*phage bubbles*). Phables models the assembly graphs of these components as a minimum flow decomposition problem using read coverage and paired-end mapping information and recovers the genomic paths of different variants. Finally, experimental results confirmed that Phables recovers complete and high-quality phage genomes, including important regions such as terminal repeats.

Decomposing assembly graphs has become a popular method to untangle genomes and recover strains from assemblies. However, much work needs to be done in terms of handling viral metagenomes and recovering phage genomes. In the future, we intend to add support for long-read assemblies from dedicated metagenomic assemblers that will enable Phables to enforce longer subpaths that will span across more sequences during the flow decomposition modelling. We also intend to extend the capabilities of Phables to recover linear phage genomes and explore the avenues for recovering high-quality eukaryotic viral genomes from metagenomes.

Data and Code Availability

All the datasets containing raw sequencing data used for this work are publicly available from their respective studies. The Lake Water dataset was downloaded from NCBI with BioProject number PRJNA756429, the Paddy Soil dataset from BioProject number PRJNA756429, the Wastewater dataset from BioProject number PRJNA434744, and the whole genome sequencing runs of the IBD data from BioProject number PRJEB7772. All the assembled data and results from all the tools are available on Zenodo at https://zenodo.org/record/7780089.

The code of Phables is freely available on GitHub under the MIT license and can be found at https://github.com/Vini2/Phables. All analyses in this study were performed using Phables v.0.2.0 with default parameters. Phables is also available as a package on bioconda at https://anaconda.org/bioconda/phables and on PyPI at https://pypi.org/project/phables/.

Author Contributions

V.M. designed the methods, developed the software, performed all analyses and wrote the paper. M.J.R. preprocessed the datasets. M.J.R. and P.D. assisted with developing the pipeline, optimising the steps and writing the paper. M.J.R., B.P., S.R.G., P.D., G.B. and L.K.I. tested the software and assisted with the data analysis. S.K.G., R.D.H. and A.L.K.H. curated data. All the authors reviewed the manuscript and provided detailed feedback. P.D., E.A.D. and R.A.E. conceived the project and wrote the paper with input from all authors.

Acknowledgements

We thank Prof Christopher Quince for his insights and suggestions regarding the development of methods. This research was undertaken with the resources and services from Flinders University's High-Performance Computing platform (“Welcome to the DeepThought HPC — DeepThought HPC Documentation” 2022).

Funding

This work is supported by the National Institutes of Health (NIH) National Institute of Diabetes and Digestive and Kidney Diseases [RC2DK116713], the Australian Research Council [DP220102915], and the Polish National Agency for Academic Exchange (NAWA) Bekker Programme [BPN/BEK/2021/1/00416 to P.D.].

Conflict of Interest

None declared.

References

Albertsen, Mads, Philip Hugenholtz, Adam Skarshewski, Kåre L. Nielsen, Gene W. Tyson, and Per H. Nielsen. 2013. “Genome Sequences of Rare, Uncultured Bacteria Obtained by Differential Coverage Binning of Multiple Metagenomes.” *Nature Biotechnology* 31 (6): 533–38.

Amgarten, Deyvid, Lucas P. P. Braga, Aline M. da Silva, and João C. Setubal. 2018. “MARVEL, a Tool for Prediction of Bacteriophage Sequences in Metagenomic Bins.” *Frontiers in*

Genetics. https://doi.org/10.3389/fgene.2018.00304.
Auslander, Noam, Ayal B. Gussow, Sean Benler, Yuri I. Wolf, and Eugene V. Koonin. 2020. "Seeker: Alignment-Free Identification of Bacteriophage Genomes by Deep Learning." *Nucleic Acids Research* 48 (21): e121.
Baaijens, Jasmijn A., Leen Stougie, and Alexander Schönhuth. 2020. "Strain-Aware Assembly of Genomes from Mixed Samples Using Flow Variation Graphs." https://doi.org/10.1101/645721.
Bankevich, Anton, Sergey Nurk, Dmitry Antipov, Alexey A. Gurevich, Mikhail Dvorkin, Alexander S. Kulikov, Valery M. Lesin, et al. 2012. "SPAdes: A New Genome Assembly Algorithm and Its Applications to Single-Cell Sequencing." *Journal of Computational Biology: A Journal of Computational Molecular Cell Biology* 19 (5): 455–77.
Belcaid, Mahdi, Anne Bergeron, and Guylaine Poisson. 2010. "Mosaic Graphs and Comparative Genomics in Phage Communities." *Journal of Computational Biology: A Journal of Computational Molecular Cell Biology* 17 (9): 1315–26.
Breitbart, Mya, Peter Salamon, Bjarne Andresen, Joseph M. Mahaffy, Anca M. Segall, David Mead, Farooq Azam, and Forest Rohwer. 2002. "Genomic Analysis of Uncultured Marine Viral Communities." *Proceedings of the National Academy of Sciences of the United States of America* 99 (22): 14250–55.
Casjens, Sherwood R., and Eddie B. Gilcrease. 2009. "Determining DNA Packaging Strategy by Analysis of the Termini of the Chromosomes in Tailed-Bacteriophage Virions." *Methods in Molecular Biology* 502: 91–111.
Comeau, André M., Graham F. Hatfull, Henry M. Krisch, Debbie Lindell, Nicholas H. Mann, and David Prangishvili. 2008. "Exploring the Prokaryotic Virosphere." *Research in Microbiology* 159 (5): 306–13.
Cook, Ryan, Nathan Brown, Tamsin Redgwell, Branko Rihtman, Megan Barnes, Martha Clokie, Dov J. Stekel, Jon Hobman, Michael A. Jones, and Andrew Millard. 2021. "INfrastructure for a PHAge REference Database: Identification of Large-Scale Biases in the Current Collection of Cultured Phage Genomes." *PHAGE (New Rochelle, N.Y.)* 2 (4): 214–23.
Delgado, Luis Fernando, and Anders F. Andersson. 2022. "Evaluating Metagenomic Assembly Approaches for Biome-Specific Gene Catalogues." *Microbiome* 10 (1): 72.
Dias, Fernando H. C., Lucia Williams, Brendan Mumey, and Alexandru I. Tomescu. 2022. "Efficient Minimum Flow Decomposition via Integer Linear Programming." *Journal of Computational Biology: A Journal of Computational Molecular Cell Biology* 29 (11): 1252–67.
Dupont, Chris L., Douglas B. Rusch, Shibu Yooseph, Mary-Jane Lombardo, R. Alexander Richter, Ruben Valas, Mark Novotny, et al. 2012. "Genomic Insights to SAR86, an Abundant and Uncultivated Marine Bacterial Lineage." *The ISME Journal* 6 (6): 1186–99.
Eddy, Sean R. 2011. "Accelerated Profile HMM Searches." *PLoS Computational Biology* 7 (10): e1002195.
Edwards, Robert A., and Forest Rohwer. 2005. "Viral Metagenomics." *Nature Reviews. Microbiology* 3 (6): 504–10.
Freire, Borja, Susana Ladra, Jose R. Paramá, and Leena Salmela. 2021. "Inference of Viral Quasispecies with a Paired de Bruijn Graph." *Bioinformatics*. https://doi.org/10.1093/bioinformatics/btaa782.
Freire, Borja, Susana Ladra, Jose R. Parama, and Leena Salmela. 2022. "ViQUF: De Novo Viral Quasispecies Reconstruction Using Unitig-Based Flow Networks." *IEEE/ACM Transactions on Computational Biology and Bioinformatics / IEEE, ACM* PP (July). https://doi.org/10.1109/TCBB.2022.3190282.
Gatter, Thomas, and Peter F. Stadler. 2019. "Ryūtō: Network-Flow Based Transcriptome Reconstruction." *BMC Bioinformatics* 20 (1): 190.
Gunasekera, Samantha, Sam Abraham, Marc Stegger, Stanley Pang, Penghao Wang, Shafi

Sahibzada, and Mark O'Dea. 2021. "Evaluating Coverage Bias in next-Generation Sequencing of Escherichia Coli." *PloS One* 16 (6): e0253440.

Hatfull, Graham F. 2008. "Bacteriophage Genomics." *Current Opinion in Microbiology* 11 (5): 447–53.

Hesse, Ryan D., Michael Roach, Emma N. Kerr, Bhavya Papudeshi, Laís F. O. Lima, Asha Z. Goodman, Lisa Hoopes, et al. 2022. "Phage Diving: An Exploration of the Carcharhinid Shark Epidermal Virome." *Viruses* 14 (9). https://doi.org/10.3390/v14091969.

Hugenholtz, P., B. M. Goebel, and N. R. Pace. 1998. "Impact of Culture-Independent Studies on the Emerging Phylogenetic View of Bacterial Diversity." *Journal of Bacteriology* 180 (18): 4765–74.

Johansen, Joachim, Damian R. Plichta, Jakob Nybo Nissen, Marie Louise Jespersen, Shiraz A. Shah, Ling Deng, Jakob Stokholm, et al. 2022. "Genome Binning of Viral Entities from Bulk Metagenomics Data." *Nature Communications* 13 (1): 965.

Jurtz, Vanessa Isabell, Julia Villarroel, Ole Lund, Mette Voldby Larsen, and Morten Nielsen. 2016. "MetaPhinder—Identifying Bacteriophage Sequences in Metagenomic Data Sets." *PLOS ONE*. https://doi.org/10.1371/journal.pone.0163111.

Kang, Dongwan D., Feng Li, Edward Kirton, Ashleigh Thomas, Rob Egan, Hong An, and Zhong Wang. 2019. "MetaBAT 2: An Adaptive Binning Algorithm for Robust and Efficient Genome Reconstruction from Metagenome Assemblies." *PeerJ* 7 (July): e7359.

Kececioglu, J. D., and E. W. Myers. 1995. "Combinatorial Algorithms for DNA Sequence Assembly." *Algorithmica*. https://doi.org/10.1007/bf01188580.

Keen, Eric C. 2015. "A Century of Phage Research: Bacteriophages and the Shaping of Modern Biology." *BioEssays: News and Reviews in Molecular, Cellular and Developmental Biology* 37 (1): 6–9.

Kieft, Kristopher, Zhichao Zhou, and Karthik Anantharaman. 2020. "VIBRANT: Automated Recovery, Annotation and Curation of Microbial Viruses, and Evaluation of Viral Community Function from Genomic Sequences." *Microbiome* 8 (1): 90.

Kolmogorov, Mikhail, Jeffrey Yuan, Yu Lin, and Pavel A. Pevzner. 2019. "Assembly of Long, Error-Prone Reads Using Repeat Graphs." *Nature Biotechnology* 37 (5): 540–46.

Lamurias, Andre, Mantas Sereika, Mads Albertsen, Katja Hose, and Thomas Dyhre Nielsen. 2022. "Metagenomic Binning with Assembly Graph Embeddings." *Bioinformatics* 38 (19): 4481–87.

Li, Dinghua, Chi-Man Liu, Ruibang Luo, Kunihiko Sadakane, and Tak-Wah Lam. 2015. "MEGAHIT: An Ultra-Fast Single-Node Solution for Large and Complex Metagenomics Assembly via Succinct de Bruijn Graph." *Bioinformatics*. https://doi.org/10.1093/bioinformatics/btv033.

Li, Heng. 2018. "Minimap2: Pairwise Alignment for Nucleotide Sequences." *Bioinformatics* 34 (18): 3094–3100.

Li, Heng, Bob Handsaker, Alec Wysoker, Tim Fennell, Jue Ruan, Nils Homer, Gabor Marth, Goncalo Abecasis, Richard Durbin, and 1000 Genome Project Data Processing Subgroup. 2009. "The Sequence Alignment/Map Format and SAMtools." *Bioinformatics* 25 (16): 2078–79.

Lima-Mendez, Gipsi, Ariane Toussaint, and Raphael Leplae. 2011. "A Modular View of the Bacteriophage Genomic Space: Identification of Host and Lifestyle Marker Modules." *Research in Microbiology* 162 (8): 737–46.

Luque, Antoni, Sean Benler, Diana Y. Lee, Colin Brown, and Simon White. 2020. "The Missing Tailed Phages: Prediction of Small Capsid Candidates." *Microorganisms* 8 (12): 1944.

Łusiak-Szelachowska, Marzanna, Beata Weber-Dąbrowska, Ewa Jończyk-Matysiak, Renata Wojciechowska, and Andrzej Górski. 2017. "Bacteriophages in the Gastrointestinal Tract and Their Implications." *Gut Pathogens* 9 (August): 44.

Mallawaarachchi, Vijini G., Anuradha S. Wickramarachchi, and Yu Lin. 2020. "GraphBin2:

Refined and Overlapped Binning of Metagenomic Contigs Using Assembly Graphs." In *20th International Workshop on Algorithms in Bioinformatics (WABI 2020)*, 21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

———. 2021. "Improving Metagenomic Binning Results with Overlapped Bins Using Assembly Graphs." *Algorithms for Molecular Biology: AMB* 16 (1): 3.

Mallawaarachchi, Vijini, and Yu Lin. 2022a. "MetaCoAG: Binning Metagenomic Contigs via Composition, Coverage and Assembly Graphs." *Research in Computational Molecular Biology: ... Annual International Conference, RECOMB ...: Proceedings. International Conference on Research in Computational Molecular Biology*, 70–85.

———. 2022b. "Accurate Binning of Metagenomic Contigs Using Composition, Coverage, and Assembly Graphs." *Journal of Computational Biology: A Journal of Computational Molecular Cell Biology* 29 (12): 1357–76.

Mallawaarachchi, Vijini, Anuradha Wickramarachchi, and Yu Lin. 2020. "GraphBin: Refined Binning of Metagenomic Contigs Using Assembly Graphs." *Bioinformatics* 36 (11): 3307–13.

McNair, Katelyn, Barbara A. Bailey, and Robert A. Edwards. 2012. "PHACTS, a Computational Approach to Classifying the Lifestyle of Phages." *Bioinformatics* 28 (5): 614–18.

Merrill, Bryan D., Andy T. Ward, Julianne H. Grose, and Sandra Hope. 2016. "Software-Based Analysis of Bacteriophage Genomes, Physical Ends, and Packaging Strategies." *BMC Genomics* 17 (1): 679.

Nayfach, Stephen, Antonio Pedro Camargo, Frederik Schulz, Emiley Eloe-Fadrosh, Simon Roux, and Nikos C. Kyrpides. 2021. "CheckV Assesses the Quality and Completeness of Metagenome-Assembled Viral Genomes." *Nature Biotechnology* 39 (5): 578–85.

Nayfach, Stephen, David Páez-Espino, Lee Call, Soo Jen Low, Hila Sberro, Natalia N. Ivanova, Amy D. Proal, et al. 2021. "Metagenomic Compendium of 189,680 DNA Viruses from the Human Gut Microbiome." *Nature Microbiology* 6 (7): 960–70.

Nissen, Jakob Nybo, Joachim Johansen, Rosa Lundbye Allesøe, Casper Kaae Sønderby, Jose Juan Almagro Armenteros, Christopher Heje Grønbech, Lars Juhl Jensen, et al. 2021. "Improved Metagenome Binning and Assembly Using Deep Variational Autoencoders." *Nature Biotechnology* 39 (5): 555–60.

Norman, Jason M., Scott A. Handley, Megan T. Baldridge, Lindsay Droit, Catherine Y. Liu, Brian C. Keller, Amal Kambal, et al. 2015. "Disease-Specific Alterations in the Enteric Virome in Inflammatory Bowel Disease." *Cell* 160 (3): 447–60.

Nurk, Sergey, Dmitry Meleshko, Anton Korobeynikov, and Pavel A. Pevzner. 2017. "metaSPAdes: A New Versatile Metagenomic Assembler." *Genome Research* 27 (5): 824–34.

Peng, Yu, Henry C. M. Leung, S. M. Yiu, and Francis Y. L. Chin. 2012. "IDBA-UD: A *de Novo* Assembler for Single-Cell and Metagenomic Sequencing Data with Highly Uneven Depth." *Bioinformatics*. https://doi.org/10.1093/bioinformatics/bts174.

Quince, Christopher, Sergey Nurk, Sebastien Raguideau, Robert James, Orkun S. Soyer, J. Kimberly Summers, Antoine Limasset, A. Murat Eren, Rayan Chikhi, and Aaron E. Darling. 2021. "STRONG: Metagenomics Strain Resolution on Assembly Graphs." *Genome Biology* 22 (1): 214.

Rho, Mina, Haixu Tang, and Yuzhen Ye. 2010. "FragGeneScan: Predicting Genes in Short and Error-Prone Reads." *Nucleic Acids Research*. https://doi.org/10.1093/nar/gkq747.

Roach, Michael J., Sarah J. Beecroft, Kathie A. Mihindukulasuriya, Leran Wang, Anne Paredes, Kara Henry-Cocks, Lais Farias Oliveira Lima, Elizabeth A. Dinsdale, Robert A. Edwards, and Scott A. Handley. 2022. "Hecatomb: An End-to-End Research Platform for Viral Metagenomics." *bioRxiv*. https://doi.org/10.1101/2022.05.15.492003.

Roach, Michael J., Katelyn McNair, Maciej Michalczyk, Sarah K. Giles, Laura K. Inglis, Evan Pargin, Jakub Barylski, Simon Roux, Przemysław Decewicz, and Robert A. Edwards. 2022.

"Philympics 2021: Prophage Predictions Perplex Programs." *F1000Research* 10 (758): 758.
Roach, Michael, Nuri Tessa Pierce-Ward, Radoslaw Suchecki, Vijini Mallawaarachchi, Bhavya Papudeshi, Scott A. Handley, C. Titus Brown, Nathan S. Watson-Haigh, and Robert A. Edwards. 2022. "Ten Simple Rules and a Template for Creating Workflows-as-Applications." *PLoS Computational Biology* 8 (12): e1010705.
Rodriguez-Valera, Francisco, Ana-Belen Martin-Cuadrado, Beltran Rodriguez-Brito, Lejla Pasić, T. Frede Thingstad, Forest Rohwer, and Alex Mira. 2009. "Explaining Microbial Population Genomics through Phage Predation." *Nature Reviews. Microbiology* 7 (11): 828–36.
Sanger, F., G. M. Air, B. G. Barrell, N. L. Brown, A. R. Coulson, C. A. Fiddes, C. A. Hutchison, P. M. Slocombe, and M. Smith. 1977. "Nucleotide Sequence of Bacteriophage Phi X174 DNA." *Nature* 265 (5596): 687–95.
Schrijver, Alexander. 1998. *Theory of Linear and Integer Programming*. John Wiley & Sons.
Shao, Mingfu, and Carl Kingsford. 2017. "Accurate Assembly of Transcripts through Phase-Preserving Graph Decomposition." *Nature Biotechnology*. https://doi.org/10.1038/nbt.4020.
Simmonds, Peter, Mike J. Adams, Mária Benkő, Mya Breitbart, J. Rodney Brister, Eric B. Carstens, Andrew J. Davison, et al. 2017. "Consensus Statement: Virus Taxonomy in the Age of Metagenomics." *Nature Reviews. Microbiology* 15 (3): 161–68.
Steinegger, Martin, and Johannes Söding. 2017. "MMseqs2 Enables Sensitive Protein Sequence Searching for the Analysis of Massive Data Sets." *Nature Biotechnology* 35 (11): 1026–28.
Terzian, Paul, Eric Olo Ndela, Clovis Galiez, Julien Lossouarn, Rubén Enrique Pérez Bucio, Robin Mom, Ariane Toussaint, Marie-Agnès Petit, and François Enault. 2021. "PHROG: Families of Prokaryotic Virus Proteins Clustered Using Remote Homology." *NAR Genomics and Bioinformatics* 3 (3): lqab067.
Tomescu, Alexandru I., Anna Kuosmanen, Romeo Rizzi, and Veli Mäkinen. 2013. "A Novel Min-Cost Flow Method for Estimating Transcript Expression with RNA-Seq." *BMC Bioinformatics* 14 Suppl 5 (Suppl 5): S15.
Tomescu, Alexandru I., and Paul Medvedev. 2016. "Safe and Complete Contig Assembly Via Omnitigs." *Lecture Notes in Computer Science*. https://doi.org/10.1007/978-3-319-31957-5_11.
Twort, F. W. 1915. "AN INVESTIGATION ON THE NATURE OF ULTRA-MICROSCOPIC VIRUSES." *The Lancet*. https://doi.org/10.1016/s0140-6736(01)20383-3.
Vatinlen, B., F. Chauvet, P. Chrétienne, and P. Mahey. 2008. "Simple Bounds and Greedy Algorithms for Decomposing a Flow into a Minimal Set of Paths." *European Journal of Operational Research*. https://doi.org/10.1016/j.ejor.2006.05.043.
"Welcome to the DeepThought HPC — DeepThought HPC Documentation." 2022. 2022. https://doi.org/10.25957/FLINDERS.HPC.DEEPTHOUGHT.
Wick, Ryan R., Mark B. Schultz, Justin Zobel, and Kathryn E. Holt. 2015. "Bandage: Interactive Visualization of de Novo Genome Assemblies." *Bioinformatics* 31 (20): 3350–52.
Xue, Hansheng, Vijini Mallawaarachchi, Yujia Zhang, Vaibhav Rajan, and Yu Lin. 2022. "RepBin: Constraint-Based Graph Representation Learning for Metagenomic Binning." *Proceedings of the AAAI Conference on Artificial Intelligence*. https://doi.org/10.1609/aaai.v36i4.20388.
Yeon-Bo Chung et al. 1990. "Bacteriophage T7 DNA Packaging: III. A 'Hairpin' End Formed on T7 Concatemers May Be an Intermediate in the Processing Reaction." *Journal of Molecular Biology* 216 (4): 939–48.
Zhang, Xing, and F. William Studier. 2004. "Multiple Roles of T7 RNA Polymerase and T7 Lysozyme during Bacteriophage T7 Infection." *Journal of Molecular Biology* 340 (4): 707–30.