

Genome analysis

# Recycler: an algorithm for detecting plasmids from *de novo* assembly graphs

Roye Rozov<sup>1</sup>, Aya Brown Kav<sup>2</sup>, David Bogumil<sup>2</sup>, Naama Shterzer<sup>2</sup>,  
Eran Halperin<sup>1,3,4</sup>, Itzhak Mizrahi<sup>2,\*</sup> and Ron Shamir<sup>1</sup>

<sup>1</sup>Blavatnik School of Computer Science, Tel-Aviv University, Tel Aviv, Israel, <sup>2</sup>The Department of Life Sciences & the National Institute for Biotechnology in the Negev, Ben-Gurion University of the Negev, Beer-Sheva, Israel, <sup>3</sup>Molecular Microbiology and Biotechnology Department, Tel-Aviv University, Tel Aviv, Israel and <sup>4</sup>International Computer Science Institute, Berkeley, CA, USA

\*To whom correspondence should be addressed.

Associate Editor: Prof. Alfonso Valencia

Received on June 25, 2016; revised on September 22, 2016; accepted on October 9, 2016

## Abstract

**Motivation:** Plasmids and other mobile elements are central contributors to microbial evolution and genome innovation. Recently, they have been found to have important roles in antibiotic resistance and in affecting production of metabolites used in industrial and agricultural applications. However, their characterization through deep sequencing remains challenging, in spite of rapid drops in cost and throughput increases for sequencing. Here, we attempt to ameliorate this situation by introducing a new circular element assembly algorithm, leveraging assembly graphs provided by a conventional *de novo* assembler and alignments of paired-end reads to assemble cyclic sequences likely to be plasmids, phages and other circular elements.

**Results:** We introduce Recycler, the first tool that can extract complete circular contigs from sequence data of isolate microbial genomes, plasmidome and metagenome sequence data. We show that Recycler greatly increases the number of true plasmids recovered relative to other approaches while remaining highly accurate. We demonstrate this trend via simulations of plasmidomes, comparisons of predictions with reference data for isolate samples, and assessments of annotation accuracy on metagenome data. In addition, we provide validation by DNA amplification of 77 plasmids predicted by Recycler from the different sequenced samples in which Recycler showed mean accuracy of 89% across all data types— isolate, microbiome and plasmidome.

**Availability and Implementation:** Recycler is available at <http://github.com/Shamir-Lab/Recycler>

**Contact:** [imizrahi@bgu.ac.il](mailto:imizrahi@bgu.ac.il)

**Supplementary information:** [Supplementary data](#) are available at *Bioinformatics* online.

## 1 Introduction

Plasmids are extra-chromosomal DNA segments carried by bacterial hosts. They are usually shorter than host chromosomes, circular and encode nonessential genes. These genes are responsible for either plasmid-specific roles such as self-replication and transfer, or context-specific roles that can be beneficial or harmful to the host depending on its environment. Along with viruses and transposable elements, plasmids are members of the group termed mobile genetic elements (Doring and Starlinger, 1984) as they transmit genes and

their selectable functions between microbial genomes. Plasmids play a central role in horizontal gene transfer (Halary *et al.*, 2009), and thus genome innovation and plasticity—fundamental forces in microbial evolution. Much interest has recently arisen for plasmid extraction and characterization, in particular because of their known roles in antibiotic resistance and in increasing metabolic outputs of agricultural or industrial byproducts. For instance, antibacterial resistance genes encoded on plasmids have long been known as a major issue for human health in clinical practice (Neu, 1992), but are also one of today's standard tools in microbiology and genetics

when used to select for specific cells (Bevan *et al.*, 1983). In order to derive plasmid sequences (which may be known or novel), one may choose from the following approaches: sequence already isolated microbes with their residing plasmids, sequence the overall microbial community of genomes (termed metagenome) from some environment or, as was recently described, sequence only the overall plasmid fraction from a given environment [termed plasmidome (Brown Kav *et al.*, 2012, 2013)]. The first technique obtains a mixture of chromosomal and plasmid DNA occurring together in a single strain. Since sequenced reads are devoted to only a few different sequenced DNA elements (the genome in question or any of its mobile elements), each is expected to be highly covered, and thus for species having low repeat content a good assembly can be achieved.

For natural environments containing many elements, often including those that are difficult to culture (Gilbert and Dupont, 2011) in a lab, metagenome assembly is attempted. This technique allows a much broader view of all taxa present and their plasmids, but is limited in that the characterization of each individual element depends on its coverage in the mixed DNA sample and the frequency of co-occurring repeats shared among different elements of the sample. Resulting assembled genomes of elements that are rare in the environment are thus often fragmented, and very high coverage (Howe *et al.*, 2014) is needed for accurately assembling them. However, assembly of metagenomes remains a highly active area of research: current assembly outputs are lacking and do not represent the true genetic capacity and synteny of genomes present in complex microbial communities. Since most of the DNA in these environments is due to host genomes, this approach currently provides only limited resolution of plasmids.

Most recently, a third technique has emerged that allows recovery of far greater numbers of plasmids. Plasmidome sequencing (Brown Kav *et al.*, 2012, 2013; Jørgensen *et al.*, 2014) allows nearly all sequencing resources to be devoted to circular DNA. Using a protocol described in (Brown Kav *et al.*, 2012), chromosomal DNA is filtered out and circular DNA segments are selectively amplified. Based on this protocol, hundreds of new plasmids were identified in the cow rumen (Brown Kav *et al.*, 2013) and rat cecum (Jørgensen *et al.*, 2014). Jørgensen *et al.*, (2014) applied the protocol introduced in Brown Kav *et al.* (2012) combined with bioinformatic validation of circularity. This post-assembly analysis resulted in a 95% PCR validation rate out of 40 randomly selected assembled contigs. This success raises the prospect of *in silico* refinement of plasmids beyond the initial assembly. Although Jørgensen *et al.*'s method was shown to have a high validation rate, its output is limited by the contiguity of the underlying assembler's contigs [in their case IDBA-UD (Peng *et al.*, 2012)], because it provides no means of combining multiple overlapping contigs to form cycles. It is a filtering process meant to identify probable circular sequences among sequences already output by the assembler. To date, no tools for plasmid assembly from short reads have been introduced to address these limitations.

In all of the above approaches plasmid assembly is hindered by several inherent characteristics derived from their mobile nature. These characteristics include their tendency to carry repetitive elements such as insertion sequences and to share genes with other plasmids and microbial genomes. In the context of *de novo* assembly, repeats cause collapse of linear sequences sharing them as subsequences. This creates ambiguity in the sense that it becomes unclear which extensions entering the repeat should be paired with those exiting it, where sequences begin and end, and whether there are unique terminal points at all as opposed to the sequence being circular. *De novo* assembly for the sake of identifying plasmids can be augmented by long-read sequencing (Conlan *et al.*, 2014; Hunt

*et al.*, 2015) because such reads may be sufficiently long to bridge repeats short reads cannot. However, this approach is primarily limited to isolates or low complexity environments. This is evident in that long reads often depend on single molecule sequencing without amplification, thus only capturing relatively abundant DNA fragments. Besides repeats, chimeric sequences also present significant challenges to assembly, in that they create false connections between sequences and thus may lead to mis-assemblies.

To overcome some of these challenges, Antipov *et al.*, (2016), introduced plasmidSPAdes, an extension of the SPAdes assembler (Bankevich *et al.* (2012) that identifies likely 'plasmid components' in isolate whole genome sequencing experiments. This method looks for long contigs in the assembly graph that are sufficiently different coverage from those of the host genome. Here, we take a different approach to improve discovery of sequenced plasmids. We similarly analyze assembly graphs, but consider all nodes instead of paring the graph around long contigs. In addition to coverage, we also incorporate paired-end read mappings and topology, only reporting cycles when there is sufficient evidence that they are physically separate entities. We also accept as input any assembly graph, making our method applicable to isolate as well as metagenome and plasmidome samples.

Our inputs are an assembly graph  $G=(V,E)$ , and the mapping of paired-end reads responsible for the assembly to its nodes. The set of nodes  $V$  are sequences having associated lengths and coverage levels, and the set of arcs  $E$  is composed of directed connections among the nodes. Arcs are the result of branch points in the underlying de Bruijn graph: a branch node has outgoing arcs to two (or more) different nodes based on overlaps, and in many cases, the assembler does not have a definite way of choosing which extension is true in order to simplify the branch into a linear path. We aim to generate a set of putative cycles that are likely to be plasmids, and assign a coverage level for each one.

After defining this problem formally below, we present an algorithm (and its implementation) designed to address it, called Recycler. Recycler leverages assembly graphs output by SPAdes to specifically enable *de novo* assembly of plasmids and other cyclic sequences likely to be physically separated from the rest of the sequences present. We show it greatly improves recovery of plasmids over naive assembly and alternative methods, namely Jørgensen's and SPAdes' built-in repeat resolution, introduced in (Prjibelski *et al.*, 2014) and performs similarly to plasmidSPAdes on isolate sample inputs. We demonstrate Recycler's performance by applying it on both simulated and real data. We find that Recycler greatly increases recall while maintaining high precision. This is established via comparisons performed on simulated plasmidomes of various sizes. We also show that Recycler can be applied for plasmid assembly on real data from a bovine rumen plasmidome and metagenome, and from two different *Escherichia coli* isolate strains. In the isolate cases, Recycler recovered most known plasmids, and predicted additional sequences that matched known mobile elements from different hosts—all of which were identical or nearly identical to known reference sequences. In all cases on real data, Recycler either matched or exceeded the proportion of outputs matching plasmid annotation, as described in Brown Kav *et al.* (2013).

## 1.1 Related work

We note plasmid assembly is a multi-assembly problem, as described in the context of RNA-Seq transcriptome assembly (Pertea *et al.*, 2015). Formulations of such problems often aim to generate a minimal set of paths that maximize agreement with observed data

(Perteau *et al.*, 2015; Tomescu *et al.*, 2013; Trapnell *et al.*, 2010). These methods usually employ network flow formulations, which admit polynomial-time algorithms for minimizing flow cost on the network; this flow corresponds to a convex function of the sum of coverage differences between observed and estimated coverage levels. However, these methods resort to heuristics in selecting a minimal set of paths to cover the entire graph, as splitting a flow into a minimal number of path and cycle components is an NP-hard problem (Hartman *et al.*, 2012).

Recycler does not aim to generate a set of paths explaining all coverage levels, and thus does not depend on a global objective function encompassing all nodes or edges. This approach is avoided because of the presence of linear paths due to either plasmids not fully covered during sequencing or bacterial host genomes housing plasmids, which may introduce noise into coverage levels observed and will not be part of the solution. Avoiding a global objective imposing parsimony on paths also allows Recycler to benefit from a polynomial time algorithm for generating ‘good’ cycles. Thus, Recycler’s approach is similar to StringTie (Perteau *et al.*, 2015), in that both repeatedly seek locally best paths or cycles and use coverage levels estimated on those to update coverage levels on the original graph, until some stopping criterion is met. We note the set of cycles desired is explicitly not minimal, as in cycle cover formulations (Gross *et al.*, 2013). For example, given a figure 8 component (Supplementary Figure S1, panel I), Recycler may represent it as two cycles separated by distinct coverage levels, where a minimal cover would use only one cycle. Instead, we wish to cover as much of the graph as possible with ‘good’ cycles.

## 2 Methods

### 2.1 Overview of recycler

The inputs to Recycler are a FASTG file representing a directed graph with vertices corresponding to non-branching sequence contigs and edges corresponding to connecting overlapping  $k$ -mers, and a BAM file of paired-end read mappings to the graph’s nodes. The graph can be viewed as a compacted de Bruijn graph starting from order  $k$  of the sequence data by contracting edges  $(u, v)$  whenever  $u$  has outdegree 1 and  $v$  has indegree 1, and the sequence contig of the new node replacing  $u$  and  $v$  is the concatenation of their sequences. Each node has a coverage value reflecting its abundance in the input sequences. We search for cycles in the graph that will correspond to plasmids. Cycle sequence length, number of vertices and coverage uniformity are factored in the selection process. We also use paired-end read mappings including mates on different nodes as a proxy for which of the nodes may have emerged from the same physical DNA fragment. This provides a means of inferring whether a candidate cycle is a plasmid or a genomic segment including repeats that lead to ambiguous cycles in the graph. Once a best cycle is selected, its latent coverage level is determined and subtracted from those of all participating nodes. Nodes whose resulting coverage values become non-positive are then removed from the graph, allowing only those with some remaining coverage the opportunity to take part in additional cycles. Hence, the whole process can be viewed as greedily ‘peeling off’ cycles from the graph. Ideally, one would like the process to end in an empty graph, in which case the input graph would be exactly the union of the cycles found. In reality, the process is stopped when quality criteria for new cycles in the remaining graph are unmet.

### 2.2 Notations and definitions

Our input is a directed graph  $G = (V, E)$ , where  $V$  is a set of linear sequences having either a branch-point or terminal  $k$ -mer at each end and no internal branch-points.  $E$  is the set of overlaps between nodes, where  $E = \{(u, v): \text{the } (k - 1)\text{-mer suffix of } u = \text{the } (k - 1)\text{-mer prefix of } v\}$ . We call a node *simple* if its indegree and outdegree are 1. A node  $v$  corresponding to sequence  $s$  of length  $l(s)$  is assigned two positive values,  $len(v)$  and  $cov(v)$ .  $len(v) = l(s) - k + 1$  is called the length of the node (the subtraction is in order avoid double-counting bases common to overlapping segments at their ends).  $cov(v)$ , its coverage, reflects the average number of times each  $k$ -mer in  $s$  appears in the input read data. The input can be produced by a short read assembly tool. We further assign a weight  $w(v) = \frac{1}{len(v)cov(v)}$  for each node  $v$ , resulting in low weight for high coverage and long nodes. Longer contigs tend to be less prone to random fluctuations in coverage, and are thus more reliable coverage indicators. For each cycle  $c$  in the graph, we assign each node a value representing its length fraction in  $c$ :  $f(c, v) = \frac{len(v)}{\sum_{v' \in c} len(v')}$ . The value  $f(c, v)$  is used to define the mean and standard deviation of weighted coverage of cycle  $c$  as  $\mu(c) = \sum_{v \in c} f(c, v)cov(v)$  and  $STD(c) = \sqrt{\sum_{v \in c} f(c, v)(cov(v) - \mu(c))^2}$ , respectively, and consequently the coefficient of variation of  $c$ ,  $CV(c) = \frac{STD(c)}{\mu(c)}$ .  $CV(c)$  is used to allow direct comparison of variation levels between cycles, independently of the magnitude of coverage of each.  $CV(c)$  is indicative of coverage uniformity along  $c$ , and plasmids are expected to have uniform coverage levels that in many cases are different from other plasmids and their hosts. Thus, cycles with low CV values are more likely to correspond to plasmids than cycles with high CV values.

### 2.3 Our approach

Intuitively, plasmids should form cycles that are distinctive from the rest of the graph and have near uniform coverage. We also expect plasmid cycles to include few nodes, as each additional node introduced for a fixed sequence length increases fragmentation and the tendency of nodes to be common to more than one path. With this in mind, we search for ‘good cycles’ in the graph that potentially correspond to plasmids. Formally, we define a good cycle as a simple cycle in the graph satisfying the following constraints:

1. Minimum path weight for some edge:  $\exists (u, v) \in c$  such that  $c \setminus (u, v)$  (the path obtained by removing  $(u, v)$  from  $c$ ) is a minimum weight path (by sum of weights  $w(v)$ ) from  $v$  to  $u$ .
2. Low coverage variation:  $CV(c) \leq \frac{\tau}{|c|}$ , where  $\tau$  is a defined threshold and  $|c|$  is the number of nodes in.
3. Concordant read mapping: For pair  $r_1, r_2$  of paired-end mates, if  $r_1$  maps to a simple node in  $c$  then  $r_2$  must also map to some node in  $c$ .
4. Sufficient sequence length:  $\sum_{v \in c} len(v) \geq L$ , where  $L$  is a defined threshold.

The first constraint is critical in order to avoid merging of two or more plasmids that are connected through a repeated region (Supplementary Figure S1, panel I). In addition, lower weight cycles correspond to longer sequence length and higher coverage nodes, and tend to have fewer nodes. Furthermore, for each edge this constraint uniquely determines at most one cycle that passes through the edge, thus avoiding consideration or enumeration of an exponential number of possible cycles. We note there are special cases allowing for cycles that visit a single node more than once; such a

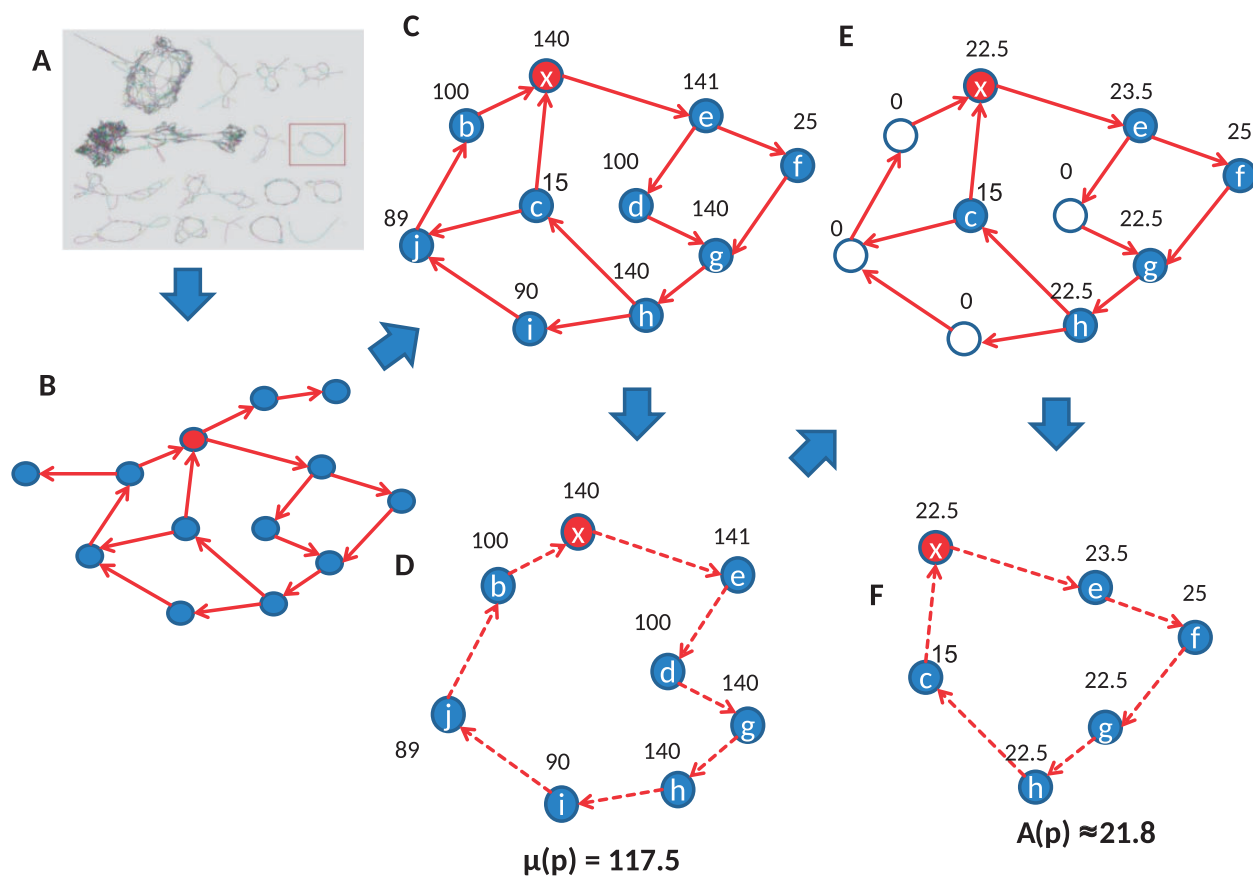
case is shown in Supplementary Figure S1, panel II. The second constraint ensures that the coverage variation is low, thus again increasing our confidence that the cycle corresponds to exactly one plasmid. Moreover, this constraint implicitly ensures high coverage cycles, since low coverage cycles tend to have higher CV value. The third constraint exploits paired-end reads. Suppose we have a read pair  $r_1, r_2$  and  $r_1$  maps to a certain node in the candidate cycle  $c$ . We expect  $r_2$  to map to the same cycle, unless  $r_1$  falls on a node that is common to  $c$  and some other path  $p$  overlapping with it. In that case  $r_2$  may map  $p$  to as well. Simple nodes are less likely to overlap with several cycles and paths, and the third constraint leverages this observation. We waive this constraint in case the coverage of  $c$  is sufficiently high, as in such cases the cycle ‘stands out’ from the background coverage. See Supplementary Material for details.

The above definition of a good cycle provides a mechanism for the identification of putative plasmids. Recycler processes each strongly connected component separately. It repeatedly finds a good cycle with minimum CV value, assigns it *latent coverage* equal to the mean cycle coverage and subtracts that coverage from the graph, creating a new *residual coverage* (Fig. 1). The weights of the vertices in the cycle are updated based on their new coverage values, and vertices whose resulting coverage values become non-positive are

removed from the graph, allowing only those with positive residual coverage the opportunity to take part in additional cycles. After each such change, cycles are recalculated the same way using the updated coverage levels. This process continues as long as new good cycles are found. To avoid examining a potentially exponential number of cycles, we consider one minimum weight cycle through each edge in the graph. The algorithm selects the cycle with the lowest CV among these minimum weight cycles and ‘peels it off’ the graph. Algorithm 1 sketches the procedure for a single component. See the Supplementary Material for additional details.

## 2.4 Complexity

Algorithm 1 presented above terminates in polynomial time. In each iteration, if any good cycles exist, one is chosen and its mean coverage is calculated. There is at least one node in the cycle with coverage smaller than the mean coverage of the cycle, which is subsequently removed from the graph. Therefore, in each iteration at least one node is removed, and the number of iterations is bounded by the number of nodes. Using Johnson’s algorithm (Johnson, 1977), the runtime of each iteration is  $O(|V|^2 \log(|V|) + |V||E|)$ . Running times are further reduced by



**Fig. 1.** Recycler work-flow. An example is shown of generating candidate cycles and peeling off cycles iteratively. For simplicity, all lengths are assumed to be equal and not shown. Here, we consider only candidate cycles that pass through vertex  $x$ , but ordinarily such candidates would be generated for each vertex in the component, and the cycle with lowest CV will be chosen and peeled off. (A) The assembly graph. (B) A single component is selected from the assembly graph (framed in A) and represented with vertices for contigs and edges for connecting  $k$ -mers. (C) The reduced component after tip removal. The numbers next to vertices are their observed contig coverage. Since vertex  $x$  has two incoming edges from vertices  $b$  and  $c$ , two candidate cycles are generated that pass through edges  $(b, x)$  and  $(c, x)$ , respectively. This is done by computing shortest paths from  $x$  to  $b$  ( $x, e, d, g, h, i, j, b, CV = 0.20$ , shown in D) and from  $x$  to  $c$  ( $x, e, d, g, h, c, CV = 0.41$ , not shown). Two successive steps of peeling cycles are shown with their respective latent coverage assignments. First, the cycle in D is peeled off because the CV calculated from initially observed coverage is lowest for this cycle. Uncolored vertices correspond to contigs with zero coverage that are removed

---

**Algorithm 1:** Finding good cycles and peeling them off each component

---

**Data:**  $G = (V, E, len, cov, w), \tau, L$

**Result:**  $\Sigma$ , the set of cycles

Compute shortest cycles passing through each edge;

for each edge  $(u, v)$  do

    Compute a minimum weight path  $p$  from  $v$  to  $u$ , if one exists;

    Compute the CV of the cycle  $(p, (u, v))$ ;

end

Return the set of cycles  $S$ ;

while  $\Sigma$  changes do

    Compute a set  $S$  of shortest cycles passing through each edge

    Consider each cycle  $c$  in  $S$  in increasing order of CV values

    if  $c$  is good and not in  $\Sigma$  then

        Add  $c$  to  $\Sigma$

        Compute the latent coverage level of  $c$

        Update the residual coverage of all cycle nodes, removing nodes with non-positive residual coverage

    else

    end

end

---

computing the strongly connected components of and working separately on each one.

## 2.5 Generating simulated plasmidomes

We simulated error-free paired-end reads from plasmids using BEAR (Johnson *et al.*, 2014), a read simulator designed to generate artificial metagenome data. To avoid introducing coverage drops at sequence ends typical of linear sequences, we modified BEAR (<https://github.com/rozovr/BEAR>) to allow sampling of reads bridging reference sequence ends, as is observed for circular sequences. Plasmid reference sequences were selected from the NCBI plasmids database and from plasmid sequences reported in (Brown Kav *et al.*, 2013), filtered to include 2760 sequences with a length range of 1–20 kbp with a mean of 6337 bp. Five datasets were created, composed of 100 bp mates (read pair ends), with insert sizes, varying from 1.25 M pairs sampled on 100 reference sequences doubling successively up to 20 M pairs sampled on 1600 sequences. Abundance levels were assigned using BEAR's low complexity option, which concentrates high abundance to few species using a power function with parameters derived from (Pignatelli and Moya, 2011): the function takes the form  $c^i d$ , where  $c=31.4$  and  $d=-1.28$ , and  $i$  is iteratively assigned values from 1 to the number of species simulated. These values are then normalized by their sum to yield a probability distribution.

## 2.6 Evaluating performance

To test recovery of the ground truth sequences by each plasmid detection program, we used the Nucmer alignment tool (Kurtz *et al.*, 2004), which is designed for efficiently comparing long nucleotide sequences such as those of whole plasmids or chromosomes. In order to simplify this process, we modified reference sequences to remove non-ACGT characters before read simulation and alignments. To avoid fragmented alignments caused by differences in start

positions, we concatenated each reference sequence to itself before mapping; this allowed identification of complete matches at the center of the concatenated contigs when they were present. Output cycles of each tested program were defined as true positives (TP) if they had 100% identity hits covering at least 80% of one of the reference sequences. False positives (FP) were any output cycles not meeting these criteria, and false negatives (FN) were reference sequences not aligned to in the output set using these criteria. Based on these conventions, precision =  $\frac{TP}{TP+FP}$  and recall =  $\frac{TP}{TP+FN}$ . We used the F1 score (Powers, 2011) to combine these measures in a manner that weighs precision and recall equally.

## 2.7 Primer design and PCR validation of plasmid contigs


The plasmidome dataset was divided into two separate subsets, including simple (single node) cycles ( $N=370$ ) and complex (multi-node) paths within the graph ( $N=50$ ). Each of these was divided into coverage bins, and selected representatives from each bin (High coverage: 60–1000x, mid–high coverage: 15–60x, mid–low coverage: 5–15x, low coverage: 1–5x) were validated by PCR. Overall, 24 simple cycles and 39 complex cycles were chosen for PCR validation. From the metagenome dataset ( $N=40$ ), all assembled plasmids were of the same coverage bin (1–5X) and 10 of them were randomly selected for validation. This was also the case for the *E. coli* E2022 isolate ( $N=4$ ) for which all plasmids were validated by PCR, aside from a recovered Phi X control sequence. Primers were designed to produce an amplification product only if their template is circular; this was achieved by directing the opposing primers towards the edge of the linear plasmid contig. PCR reactions were carried out using Advantage GC Genomic LA PCR Polymerase (Clontech) according to the manufacturer's instructions. The PCR reactions were as follows: 1.5  $\mu$ l Advantage buffer (10 $\times$ ), 0.6  $\mu$ l of each primer (5 mM), 0.15  $\mu$ l Ex Advantage GC Genomic LA DNA Polymerase, 100 ng of template DNA, 1.5  $\mu$ l of dNTPs (10 mM) and DDW was added to a final volume of 25  $\mu$ l. All PCR reactions were carried out in a Sensoquest thermocycler (Gottingen, Germany).

## 3 Results

We first simulated plasmidomes using known references. We used these data sets to assess Recycler's precision and recall (along with those of alternative methods) by comparing predictions against the ground truth known by the simulation design. We also tested Recycler on real data from two *E. coli* isolates, and both a cow rumen metagenome and plasmidome (Brown Kav *et al.*, 2013). For the bacterial isolates that have been sequenced, predicted plasmids were compared against the reference sequences directly. Since no references are available for metagenome and plasmidome data, we evaluated the accuracy by PCR validation (Jørgensen *et al.*, 2014) and by measuring the proportion of predicted plasmids having proper annotation as done in (Brown Kav *et al.*, 2013). Recycler's inputs were assembly graphs generated by SPAdes version 3.6.2 (Bankevich *et al.*, 2012), and alignments generated by BWA version 0.7.5 (Li and Durbin, 2009).

### 3.1 Simulated plasmidomes

We simulated paired-end reads from known plasmids, and created five datasets of 100, 200, 400, 8000 and 1600 plasmids. Plasmid abundance was distributed so that few plasmids have high abundance. Dataset sizes were 1.25, 2.5, 5, 10 and 20 M pairs, respectively (see Methods for details). Each such dataset was assembled with SPAdes and subsequently its output contigs and assembly



**Fig. 2** Methods performance on simulated data. Results are shown for SPAdes without repeat resolution (RR), SPAdes with repeat resolution, the method of Jørgensen *et al.*, and Recycler. The contigs of SPAdes before RR were used as input for the three other methods. Recycler also relied on the graph produced at this stage. F1 score calculation is described in the main text. The x axis shows the number of simulated reference sequences in each case

graphs were used as inputs to the tested methods. Recycler was compared with SPAdes with and without repeat resolution (RR), and with a simplified version of Jørgensen's method (described in the Appendix). We used SPAdes' outputs before the repeat resolution stage as inputs to Recycler and to Jørgensen's method, as we found that contigs have greater precision before RR when compared to reference sequences (as shown in Supplementary Table S1). The mapping results are presented in Supplementary Table S1 and Figure 2.

As expected, recall generally decreased as the number of simulated plasmids increased. This was common to all tested methods. In general, we found that Recycler generated more predictions than other methods, leading it to have higher recall than alternative approaches while maintaining high (~90%) precision. The net performance effect is shown in Figure 2 and Supplementary Table S1 in the supplement: Recycler maintains the lead in all cases with 5–14% advantage in both F1 and fraction of true positives. We also found that the number of additional Recycler true positives over those provided by SPAdes generally increased with higher complexity; this culminated in Recycler adding 62 (13%) true positives to SPAdes' output on the 1600 plasmid set (523 versus 461).

To further characterize Recycler's performance, we categorized its predictions in terms of mean total cycle length, number of segments in the cycle (steps), cycle coverage and CV value calculated at the stage the cycle was removed. For each category, values were subdivided into five ranges. In Supplementary Figure S2, we show the precision values and the relative proportions of counts in the specified ranges. Based on this stratification, it can be seen that Recycler shows little dependence on mean coverage or length, but does often preclude candidate cycles that have high CV values or number of steps. This is reflected in the sharp drop-off in the plots as the number of steps or the CV grows.

### 3.2 Real data

All of Recycler's results on real data were subjected to quantification of annotation results as described in (Brown Kav *et al.*, 2013) and compared against cycles present in the output produced by SPAdes. These results are detailed below and a summary of them can be found in Supplementary Table S2 in the Appendix.

#### 3.2.1 Circular integrity of assembled plasmids

A total of 77 sequences were selected for PCR validation by sampling from the different data types as described in the Section 2.7 above. Overall, 89% of the 77 chosen plasmids were validated by PCR as circular DNA molecules. The predicted plasmids from the different samples did not differ in the success rate of circular validation. As coverage has a key role in *de novo* assembly and Recycler's performance, we wished to measure whether the integrity of assembled plasmids would be affected by varying mean *k*-mer coverage. To this end, we validated circularity of plasmids of different coverage levels ranging from 1x to 1000x divided into bins. As can be seen in Figure 3, there was a slightly lower success rate for the lower coverage plasmids. However, coverage and validation rate were not found to be significantly correlated. Additionally, the high number of predicted plasmids in the plasmidome data set allowed us to measure the effect of the complexity of the path in the graph on the integrity of the plasmids. When more edges are involved in a cycle, it is more complex, and the chance of noise in coverage levels and errors in sequence increases. Thus, we divided this dataset into two bins according to path length on the graph: simple: single node (self-edge) paths, complex: two nodes or more. These two bins did not show difference in their validation rate, further stressing Recycler's strength in extracting plasmids from complex paths.

#### 3.2.2 *E. coli* isolate data

We ran Recycler on two *E. Coli* strains: JJ1886, downloaded from <http://www.ebi.ac.uk/ena/data/view/SRX321704>, and E2022, sequenced locally. Annotation for plasmids found in both strains was provided in (Lanza *et al.*, 2014); comparisons against Recycler outputs with this annotation are reported in Supplementary Tables S3 and S4. Of the five plasmids known for JJ1886, Recycler output four complete matches (100% identity over 100% length) having lengths 55.9, 5.6, 5.2 and 1.6 kbp. It also output three additional sequences which completely matched previously reported plasmids: two are known to be present in *S. Aureus*, and one in *S. Chromogenes*. Further tests will be needed in order to validate whether these additional hits are truly present in the sequenced sample, and furthermore, whether they are stable residents of the tested hosts or were present as a result of contamination. When tested on E2022, Recycler performed similarly, recalling most of its known plasmids and outputting a few additional cycles that were complete or near complete matches to known plasmids and one phage. These results are also presented in Supplementary Table S2. In summary, all reported isolate hits represent highly accurate matches to known mobile elements, and most known plasmids for these strains were recovered. In both cases, Recycler missed the longest known reference plasmids; it remains to be seen whether this is due to Recycler's use of a shortest path formulation, lack of significant coverage difference between these plasmids and the host genome, or other factors.

#### 3.2.3 Plasmidome data

A bovine rumen plasmidome sample was prepared as described in (Brown Kav *et al.*, 2013). This data consisted of 5.1 M paired-end 101 bp reads (trimmed to varied sizes for the sake of adapter removal) with an expected insert size of 500 bp [data available upon request]. Recycler output 420 cycles when provided this data. According to ORF prediction performed as in (Brown Kav *et al.*, 2013), 314 of the 420 had significant annotation hits. 96% of those matching annotations either matched plasmid annotations or aligned with plasmids reported in (Jørgensen *et al.*, 2014). Thus, a majority are likely to be plasmids.

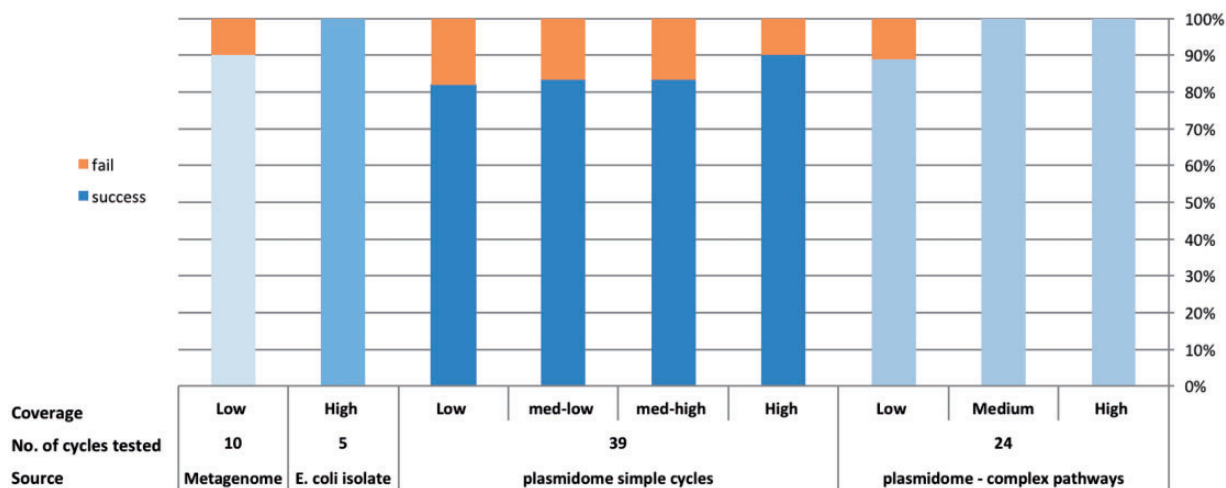


Fig. 3. PCR based validation of Recycler's plasmid predictions. High coverage: 60–1000x, med-high:15–60x, med-low: 5–15x, low: 1–5x

### 3.2.4 Metagenome data

Metagenome data was derived from the rumen of a different cow residing in the same stable as the cow used to derive the plasmidome data. This data consisted of 7.5 M paired end 150 bp reads with expected insert size of 500 bp [data available upon request]. Recycler produced 40 cycles when run on this data. According to ORF prediction, 37 of the 40 had significant annotation hits. About 35% of those matching annotations either matched plasmid annotations or aligned with plasmids reported in Jørgensen *et al.*, (2014). The proportion of reported cycles matching known plasmid annotations was slightly higher than for simple cycles output by SPAdes (33%). Overall, this test reflects the trend seen elsewhere (Howe *et al.*, 2014) of weak annotation results emerging from metagenome assembly of highly diverse environmental samples.

### 3.2.5 Comparison with PlasmidSPAdes

Recently, a version of SPAdes tailored for seeking plasmids in isolates, called PlasmidSPAdes, was introduced (Antipov *et al.*, 2016). Unlike Recycler, it does not explicitly seek cycles but removes long edges in the de Bruijn graphs and looks for contigs with coverage significantly different from the mean coverage of the read data. The rationale is that for isolates the coverage distribution is dominated by the host bacterium reads, and the reads of plasmids can be detected as outliers in that distribution. This assumption does not fit plasmidome or metagenome data. PlasmidSPAdes' output is a set of components, each containing a set of contigs with similar mean coverage that putatively originate from the same plasmid. We ran PlasmidSPAdes (packaged with SPAdes 3.80) on the two *E. coli* datasets described above, and compared the results with Recycler's (Supplementary Tables S5 and S6). For E2022, four out of the seven components reported by PlasmidSPAdes matched Recycler's outputs; the shortest of these was among the PCR validated sequences not present in the reference set. Of the three not matching, two seem to have chromosomal origin based on a BLAST search performed on the longest contigs in these components, and the fact that these components had largely tree-like structure: less than half of the component's total length was included in a cycle. Recycler reported one cycle of length 2.1 kb missed by PlasmidSPAdes that was in the reference set. Neither tool succeeded in recovering the longest two plasmids in the reference set.

For JJ1886, three out of the nine components reported matched Recycler's. Of the other six, five likely have chromosomal origin as

assessed by the same criteria used for E2022, and one matched a likely plasmid. However, four of these five aligned best with the genome of *S. Aureus*. Recycler reported three additional short sequences between 1.6 and 2.4 kb, each of which had high scoring BLAST hits to plasmids in *S. Aureus* or *S. Chromogenes*. As some of the plasmids reported by both tools also matched *S. Aureus* origin, it is possible that the JJ1886 sample contained a mixture of both cell types. We note that such a mixture could mislead PlasmidSPAdes' estimates of coverage variation, thus allowing large chromosomal fragments to survive filtration.

Overall, aside from the *S. Aureus* sequences observed, the two tools performed similarly on isolate data. This is consistent with the comparison presented in (Antipov *et al.*, 2016). In addition, Recycler can process metagenome and plasmidome graphs, while PlasmidSPAdes can find non-circular plasmids. The two methods primarily differ (when processing isolate data) in what they report for difficult cases involving repeats that are either long or shared by many paths. When Recycler cannot derive a unique circular sequence from a graph component, the component is not included in the output. For PlasmidSPAdes, such components are reported as groups of contigs. In either case, more information (such as long reads) would be needed in order to properly resolve these cases.

## 4 Discussion

In this article, we describe Recycler, a new algorithm and the first tool available for identification of plasmids from short read-length deep sequencing data. We demonstrate that Recycler discovers plasmids that remain fragmented after *de novo* assembly. We have adapted the approach of choosing among likely enumerated paths using coverage and length properties (often applied in transcriptome assembly (Perteau *et al.*, 2015; Tomescu *et al.*, 2013; Trapnell *et al.*, 2010) for extracting a specific but common inhabitant of metagenomes. We showed that many more real plasmids can be found by only generating likely cycles on the assembly graph versus alternative methods. We validated this approach on both real and simulated data.

Recycler displays high recall and precision on simulated plasmidomes, and we have developed a means of separating real plasmids from cycles due to repeats in isolate data. As we have noted, coverage can be very useful for the latter, but the assumption that coverage will always differ significantly between plasmids and their host

genome does not hold universally. It is worth noting that as new plasmids are identified and their common sequence motifs are observed, both reference-based identification and a priori trained prediction of plasmid features can be improved and harnessed for supplementing identification based on coverage and length features alone. We aim to investigate how such knowledge can be leveraged for increased precision without sacrificing recall.

Furthermore, while Recycler's peeling of lowest CV paths at each step has the advantage of providing a deterministic rule to decide which cycles should be peeled next, this process is heuristic. Better accounting of the uncertainty in observed coverage levels and in the algorithm's dependence on the order of peeling may be obtained by randomizing or repeating parts of the process multiple times. For example, instead of always peeling one best cycle, a random subset of all good cycles may be peeled at once. Repeating this process multiple times and reporting only cycles that persist in a majority of runs may improve both sensitivity and precision.

Further investigation will be needed to assess how plasmids can be extracted from environmental samples, in spite of the limitations now hampering metagenome assembly. This is currently challenging, as diverse genomes require very high coverage for rare species to be captured, but such high coverage data demand computational resources beyond reach of most investigators. While new techniques have aimed to address this problem (Cleary et al., 2015; Howe et al., 2014), they have yet to see widespread use, and work best when paired with multiple samples to allow for species separation by co-abundance signatures. Along with addressing these concerns, it remains to be seen whether a mixed approach of pre-screening environmental samples for plasmids and computationally filtering them out may benefit metagenome graph simplification.

## Acknowledgements

RR wishes to thank Kobi Perl and David Pellow for helpful comments given in the preparation of the manuscript.

## Funding

This work was supported in part by the Israel Science Foundation [grants no. 1425/13 to EH, 317/13 to RS, and 1313/13 to IM], and the Israel Science Foundation-National Natural Science Foundation of China joint program 2015-18 to RS. Additional support was provided by the European Research Council under the European Union's Horizon 2020 research and innovation program [grant agreement No 640384 to IM] and the Israeli Center of Research Excellence (I-CORE), Gene Regulation in Complex Human Disease, Center No [Grant 41/11 to RS]. RR was supported in part by a fellowship from the Edmond J. Safra Center for Bioinformatics at Tel Aviv University, an IBM PhD fellowship, and by the Center for Absorption in Science, the Israel Ministry of Immigrant Absorption. EH is a Faculty Fellow of the Edmond J. Safra Center for Bioinformatics at Tel Aviv University.

*Conflict of Interest:* none declared.

## References

Antipov, D. et al. (2016). plasmidSPAdes: Assembling Plasmids from Whole Genome Sequencing Data. Technical report.

- Bankevich, A. et al. (2012) SPAdes: a new genome assembly algorithm and its applications to single-cell sequencing. *J. Comput. Biol.*, **19**, 455–477.
- Bevan, M.W., Flavell, R.B. and Chilton, M.D. (1983) A chimaeric antibiotic resistance gene as a selectable marker for plant cell transformation. *Nature*, **304**, 184–187.
- Brown Kav, A. et al. (2012) Insights into the bovine rumen plasmidome. *Proc. Natl. Acad. Sci. USA*, **109**, 5452–5457.
- Brown Kav, A. et al. (2013) A method for purifying high quality and high yield plasmid DNA for metagenomic and deep sequencing approaches. *J. Microbiol. Methods*, **95**, 272–279.
- Cleary, B. et al. (2015) Detection of low-abundance bacterial strains in metagenomic datasets by eigengene partitioning. *Nat. Biotechnol.*, **33**, 1053–1060.
- Conlan, S. et al. (2014) Single-molecule sequencing to track plasmid diversity of hospital-associated carbapenemase-producing Enterobacteriaceae. *Sci. Translat. Med.*, **6**, 254ra126.
- Doring, H. and Starlinger, P. (1984) Barbara McClintock's controlling elements: now at the DNA level. *Cell*, **39**, 253–259.
- Gilbert, J.A. and Dupont, C.L. (2011) Microbial metagenomics: beyond the genome. *Annu. Rev. Mar. Sci.*, **3**, 347–371.
- Gross, J.L. et al. (2013) *Handbook of Graph Theory, 2nd edn.* Chapman & Hall/CRC, Boca Raton, FL.
- Halary, S. et al. (2009) Network analyses structure genetic diversity in independent genetic worlds. *Proc. Natl. Acad. Sci. USA*, **107**, 127–132.
- Hartman, T. et al. (2012) How to split a flow? In: *2012 Proceedings IEEE INFOCOM*, pp. 828–836..
- Howe, A.C. et al. (2014) Tackling soil diversity with the assembly of large, complex metagenomes. *Proc. Natl. Acad. Sci. USA*, **111**, 4904–4909.
- Hunt, M. et al. (2015) Circlator: automated circularization of genome assemblies using long sequencing reads. *Technical Report*.
- Johnson, D.B. (1977) Efficient algorithms for shortest paths in sparse networks. *J. ACM*, **24**, 1–13.
- Johnson, S. et al. (2014) A better sequence-read simulator program for metagenomics. *BMC Bioinformatics*, **15 Suppl 9(Suppl 9)**, S14.
- Jørgensen, T.S. et al. (2014) Hundreds of circular novel plasmids and DNA elements identified in a rat cecum metamobilome. *PLoS One*, **9**, e87924.
- Kurtz, S. et al. (2004) Versatile and open software for comparing large genomes. *Genome Biol.*, **5**, R12.
- Lanza, V.F. et al. (2014) Plasmid flux in *Escherichia Coli* ST131 sublineages, analyzed by plasmid constellation network (PLACNET), a new method for plasmid reconstruction from whole genome sequences. *PLoS Genet.*, **10**, e1004766.
- Li, H. and Durbin, R. (2009) Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics*, **25**, 1754–1760.
- Neu, H.C. (1992) The crisis in antibiotic resistance. *Science*, **257**, 1064–1073.
- Peng, Y. et al. (2012) IDBA-UD: a *de novo* assembler for single-cell and metagenomic sequencing data with highly uneven depth. *Bioinformatics*, **28**, 1420–1428.
- Pertea, M. et al. (2015) StringTie enables improved reconstruction of a transcriptome from RNA-seq reads. *Nat. Biotechnol.*, **33**, 290–295.
- Pignatelli, M. and Moya, A. (2011) Evaluating the fidelity of *de novo* short read metagenomic assembly using simulated data. *PLoS One*, **6**, e19984.
- Powers, D.M. (2011). Evaluation: from Precision, Recall and F-measure to ROC, Informedness, Markedness and Correlation. *J. Mach. Learn. Technol.*, **2**, 37–63.
- Prijbelski, A.D. et al. (2014) ExSPAnDer: a universal repeat resolver for DNA fragment assembly. *Bioinformatics*, **30**, i293–i301.
- Tomescu, A.I. et al. (2013) A novel min-cost flow method for estimating transcript expression with RNA-Seq. *BMC Bioinformatics*, **14 Suppl 5**, S15.
- Trapnell, C. et al. (2010) Transcript assembly and quantification by RNA-Seq reveals unannotated transcripts and isoform switching during cell differentiation. *Nat. Biotechnol.*, **28**, 511–515.