

PROCEEDINGS

Open Access

Ensemble analysis of adaptive compressed genome sequencing strategies

Zeinab Taghavi

From RECOMB-Seq: Fourth Annual RECOMB Satellite Workshop on Massively Parallel Sequencing
Pittsburgh, PA, USA. 31 March - 05 April 2014

Abstract

Background: Acquiring genomes at single-cell resolution has many applications such as in the study of microbiota. However, deep sequencing and assembly of all of millions of cells in a sample is prohibitively costly. A property that can come to rescue is that deep sequencing of every cell should not be necessary to capture all distinct genomes, as the majority of cells are biological replicates. Biologically important samples are often sparse in that sense. In this paper, we propose an adaptive compressed method, also known as distilled sensing, to capture all distinct genomes in a sparse microbial community with reduced sequencing effort. As opposed to group testing in which the number of distinct events is often constant and sparsity is equivalent to rarity of an event, sparsity in our case means scarcity of distinct events in comparison to the data size. Previously, we introduced the problem and proposed a distilled sensing solution based on the breadth first search strategy. We simulated the whole process which constrained our ability to study the behavior of the algorithm for the entire ensemble due to its computational intensity.

Results: In this paper, we modify our previous breadth first search strategy and introduce the depth first search strategy. Instead of simulating the entire process, which is intractable for a large number of experiments, we provide a dynamic programming algorithm to analyze the behavior of the method for the entire ensemble. The ensemble analysis algorithm recursively calculates the probability of capturing every distinct genome and also the expected total sequenced nucleotides for a given population profile. Our results suggest that the expected total sequenced nucleotides grows proportional to log of the number of cells and proportional linearly with the number of distinct genomes. The probability of missing a genome depends on its abundance and the ratio of its size over the maximum genome size in the sample. The modified resource allocation method accommodates a parameter to control that probability.

Availability: The squeezamblor 2.0 C++ source code is available at <http://sourceforge.net/projects/hyda/>. The ensemble analysis MATLAB code is available at <http://sourceforge.net/projects/distilled-sequencing/>.

Introduction

Progress in DNA amplification techniques [1] and high throughput cell cultivation methods [2,3] allow capturing of genomes at single-cell resolution. However, deep sequencing and assembly of all of the cells in a sample is prohibitively costly since there are millions sometimes billions of cells. The good news is that, to capture all distinct genomes, deep sequencing of every cell should not be

necessary as the majority of cells are biological replicates. For instance, the number of detected distinct species in the human gut was estimated to be in the order of 1,000, while the number of microbial cells in a human body, most of which reside in the gut, is in the order of 100 trillion [4]. We call this effect the *sparsity* of distinct genomes in a sizeable microbial population. Biologically important samples are often sparse in that sense. We use sparsity to capture all of the genomes in a sample.

During the last decade, the rich field of compressed sensing in non-adaptive [5-7] and adaptive (distilled sensing

Correspondence: ztaghavi@mail.colostate.edu
Computer Science Department, Colorado State University, 346 Computer
Science Building, Fort Collins, CO 80523, USA

and refinements) forms [8,9] has been developed to reduce the cost of sampling and reconstruction of sparse signals [10,11]. In the general form of the problem, both adaptive and non-adaptive methods reduce the number of sensing in comparison to the non-adaptive naive sensing, and make it proportional to the number of distinct events times log of the data size [12]. However while adaptive compressed methods may seem more cumbersome than their non-adaptive counterparts, adaptive methods often improve detection and estimation performance [12].

Our problem is an instance of a larger class of problems called *element distinctness*, which is a popular problem in massive data analysis with numerous applications and different variants including (i) finding if there are duplicates in a list, (ii) calculating the number of distinct elements (support size), and (iii) estimating the distribution of distinct element populations [13]. In the vast majority of element distinctness problems, the complexity of deciding if two elements are identical is of $O(\text{element size})$. We distinguish between different classes of the problem based on the size of an element in comparison with the size of the entire population.

In some of the problems such as estimation of the number of distinct words an author knows (e.g. Shakespeare) [14], the size of an element is very small in comparison with the size of the problem. In some others, hash functions are used to reduce the element size. Such variants of the problem have been investigated deeply to find optimal algorithms both in time and space. For data stream analysis, if n is the size of the language of elements, then the space complexity of optimal probabilistic $(1 \pm \epsilon)$ -approximation algorithm is $O(\epsilon^{-2} + \log n)$ and its time complexity is $O(n)$ [13]. In such algorithms, each element is either completely sensed or not sensed at all, i.e., no partial sensing of an element. However, there are important variants of the problem in which each element is complex, such as the case of whole genome sequences. The contributions of this paper and its predecessor [15] is (i) introduction of this other class of problems with large element size with respect to the sample size, and (ii) the first adaptive compressed method to the best of our knowledge to solve an instance of this problem in the form of finding all distinct genomes with reduced sequencing effort in a sparse microbial sample.

Assume that the distinction of two cells is based on the differences between their genomes. Therefore, the complexity of pairwise distinction is a function of the lengths of the DNA sequences, each in the order of 1,000,000 – 10,000,000 base-pairs for a bacterial cell and 3,300,000,000 base-pairs for a human cell, with an average size of m . A sample contains n cells, for instance 10,000,000 cells, where n and m are in the same order. In this problem, there are two types of cost: (i) wet-lab cost related to

sequencing, i.e., reading the DNA sequence digitally, and (ii) computational cost of genome assembly and comparison, i.e., digital reconstruction of the whole genome sequences from the sequencer output. The output reads of a sequencer are short randomly sampled subsequences of the genomic sequence which cover the genome multiple times. The number of reads that contain a genomic location is called the coverage. In the assembly, the reads are concatenated to reconstruct the whole genome. Sequence assembly is a challenging task due to sequencing errors and repetitive elements. To compare a number of sequenced read data sets, a co-assembly software such as HyDA [16] is used. The output of HyDA provides us with measures to compare the extent of similarity between the underlying genomes from which the read data sets are derived.

Wet-lab cost includes the monetary cost which is linearly proportional to the total number of base-pairs sequenced. If m is the average genome size and c is the necessary coverage, then the cost is $O(nmc)$ for the exhaustive sequencing of all cells. Computational cost includes the space complexity and time complexity of assembly and comparison. If the assembly is done using the de Bruijn graph [17], the time complexity is $O(nmc \log m)$ and space complexity is $O(nm)$. For instance, a typical real-world scenario involves $c = 20$, $m = 5,000,000$ bps, and $n = 10,000,000$ for which the exhaustive wet-lab, time, and space cost complexities would be respectively $O(10^{15})$, $O(10^{16})$, and $O(10^{13})$. The exhaustive approach is not tractable even for a small population. Hence, sublinear algorithms are needed to solve the problem.

We propose an adaptive compressed method, also known as distilled sensing. Our ultimate goal is to reduce the nm factor in wet-lab, time, and space complexities to $sm \log n$ in which s is the number of distinct genomes in the community. We cannot use the algorithms and analyses given for the classical compressive sensing approach since our sparsity is *unordered* set sparsity rather than ordered sparsity by time or space. As opposed to group testing in which the number of distinct events is often constant and sparsity is equivalent to rarity of an event, sparsity in our case means scarcity of distinct events in comparison to the data size. It is also important to note that we do not have positional access (a.k.a. random access in the computer science literature) to the DNA sequence, which limits the use of many dimensionality reduction techniques [18].

We previously defined the problem and proposed a distilled sensing solution based on the breadth first search strategy [15]. To evaluate the performance of our algorithm, we simulated the whole process including genome amplification by MDA [19], sequencing by Illumina, (co-

assembly by HyDA [16], and comparison. We proposed an adaptive resource allocation method to determine the amount of sampling of each genome in each round, which is related to the one proposed by [8,9]. Due to the computational intensity of each of those processes, we were able to demonstrate the power of our approach for a few instances of the problem, but the behavior of the algorithm for the entire ensemble is yet to be studied.

In this paper, we give a new algorithm based on the depth first search strategy and modify our previous breadth first search resource allocation and set selection. Since simulating the entire process is time-consuming, we provide a dynamic programming algorithm to analyze the behavior of the method over the entire ensemble. Our algorithm recursively calculates the probability of capturing every distinct genome and also the expected total sequenced nucleotides for a given population profile. It is important to note that even though the population is known for the ensemble analysis algorithm, the actual sensing algorithm works without that knowledge. That is our sensing algorithm can be applied to any population, even without knowing the profile. To have a clear view of the effect of each parameter on the expected cost, we assume that our model is error free at this stage. The results in this paper may lead to theoretical solutions and analysis with more complete model assumptions in the near future.

Method

Our method consists of two parts: (i) wet-lab process, and (ii) computational process. On the wet-lab side, we are assuming to have a high throughput device which is capable of isolating each cell in the sample, cultivate it, then extract the DNAs of each cultivated cell and amplify them. This device should also be capable of sampling customized amount from selected amplified DNAs, pool them, and prepare them for sequencing. If we would like to sequence more than one pool of samples in the same run, the device should uniquely barcode each pool before sending the samples for sequencing. Although there is currently no such device, one can envision automated microfluidic devices in near future based on the technologies already developed for separation, cultivation, DNA extraction, amplification, and barcoding [2,3].

The output of sequencing is a library of reads which will be demultiplexed based on the barcodes. Therefore, for each pool of sampled amplicons which is sent for sequencing, a read data set is obtained. All the read data sets at each round are co-assembled (with HyDA [16]). In the co-assembler, to each read data set a unique color is assigned. All the colors are assembled on a single de Bruijn graph. The output is a list of contigs and their colored average coverages. This provides us with a measure to compare

the similarities between the assembly of different colors. Based on those similarities, we decide if any two assemblies could potentially be from the same genome.

In the naive exhaustive approach, each isolated cell is sampled and deeply sequenced. Based on the similarity measures provided by the co-assembler, distinct genomes are then identified. In the adaptive method, at each round a number of collections of cells are selected. For each collection, the amount to be sampled from each cell is computed based on the analysis in the previous rounds. The output read data sets are analysed and the next round of sampling is calculated. We describe the details of the sampling collections and size in this section. First, let's clarify the assumptions for our model.

Model assumptions

The definition of distinct genomes may vary in different applications. We, instead of phenotypic notions like species or strains, use a quantifiable genomic measure to determine the distinction of genomes. We define two genome sequences to be *distinct* if the ratio of their differences over the whole genome size is above a threshold, called τ . That threshold is input by the user and controls a trade-off between sensitivity and specificity [15].

Let $C = \{C_1, C_2, \dots, C_n\}$ comprise the input community of cells. As described earlier, we are given a device that can sense each cell C_i partially at random, and the cost of a sensing is proportional to the sensing size, i.e., the number of nucleotides sampled. As the sensing size increases, the reconstructed genome of C_i after the assembly converges to completion. To introduce appropriate notations, let $I \subseteq C$ be a subset of the community. Let be the sensing of the aggregated cells in I , which is the superposition of all sensing taken from the cells in I , i.e., the aggregated read data set or equivalently the resulting assembly. The key observation is that if there are enough replicates of a particular distinct genome in I , then that distinct genome can be completely captured from the superposition of partial sensing of the replicates provided that the partial sensing are random and unbiased.

Comparison of assemblies of two sets

Let $I_1, I_2 \subseteq C$ be two subcollections of the input community, and A_1, A_2 the corresponding aggregate sensing. If all of the distinct cells represented in I_1 are also represented in I_2 , then we say that A_1 *subsumed* by A_2 ($A_1 \leq A_2$). In ideal world with no errors and genome variations, A_1 is called subsumed in A_2 if A_1 is a subset of A_2 . However, in real world, while two genomes are considered similar (from the same type), they may have some variations like single nucleotide polymorphisms. In addition, errors, noise, and contaminations in sequencing and assembly make the situation harder to handle just by pure

mathematical subset definition. To address this issue, the subset definition is relaxed to ignore those differences between two assemblies that are less than a threshold, τ . Therefore, subsumption [15] is defined as follows

$$A_1 \leq A_2 \text{ iff } 0 \leq D_\tau(A_1, A_2). \quad (1)$$

In this equation, $D_\tau(A_1, A_2)$ quantifies the differences in assembly of A_1 with respect to A_2 which is more than τ and is defined as follows

$$D_\tau(A_1, A_2) = \tau - \frac{||A_1 \setminus A_2||}{||A_1||}, \quad (2)$$

in which $A_1 \setminus A_2 = \{b \in A_1 | b \notin A_2\}$, $|| \cdot ||$ denotes the total assembly size. In other words, τ is the maximum differences tolerated between two genomes which are considered similar. Parameter τ is user defined and τ gives a trade-off between specificity and sensitivity of the algorithm to distinguish between two distinct genomes [15].

Search strategies

Our algorithm aims to assemble all of the distinct genomes represented in C and identify at least one cell per distinct genome. The objective is to minimize the total number of bases required to be sequenced. To reach this goal a search tree is created and explored iteratively to find the leaves which are the sequenced and assembled species. In the first iteration, the set of deeply sequenced and completely assembled distinct genomes, I , and its aggregated sensing, A , is empty. The algorithm divides the n cells C_1, \dots, C_n into two sets $I_{1,1} = \{C_1, \dots, C_{\lfloor n/2 \rfloor}\}$ and $I_{1,2} = \{C_{\lfloor n/2 \rfloor + 1}, \dots, C_n\}$. Denote $\bar{I}^1 = \{I_{1,1}, I_{1,2}\}$. In each iteration i , $I_{i,j}$'s are subsets of C and are chosen based on the results in the iteration $i - 1$. The search tree of $I_{i,j}$ to find leaves can be traversed by different methods. Here we choose two methods, breadth first search (BFS) and depth first search (DFS). In the BFS strategy, in each iteration i , all $I_{i,j}$'s are explored at the same time, while in DFS, nodes ($I_{i,j}$'s) are explored sequentially in time and analyzed one after the other.

In the recursive call on $\bar{I}^i = \{I_{i,1}, \dots, I_{i,m_i}\}$, the set of cells is sensed according to the resource allocation policy. Then, the aggregated sensing, $A_{i,j}$, for each $I_{i,j}$ is obtained by sequencing and assembly. Those $I_{i,j}$'s that contain a single cell, i.e., $|I_{i,j}| = 1$, are leaves, and if they are fully assembled, they will be added to the list of deeply sequenced cells. In other words, if the corresponding assembly is reliable, i.e., $c_{i,j} \geq M_b$, for a given constant M_b , $I_{i,j}$ will be popped from \bar{I}^i and pushed to I . In addition, $A_{i,j}$ will be added to A .

For the BFS search strategy to find the optimum path to continue, a subset of \bar{I}^i with minimum number of cells is

chosen that covers all of the assembly. In other words, the minimum assembly-set cover $\bar{I}_{cover} \subseteq \bar{I}^i$ with minimum number of cells is found for which $\bar{A} \cup A$ is subsumed in $\bar{A}_{cover} \cup A$, i.e.,

$$\begin{aligned} D_\tau(\bar{A} \cup A, \bar{A}_{cover} \cup A) &= \tau - \frac{||(\bar{A} \cup A) \setminus (\bar{A}_{cover} \cup A)||}{||\bar{A} \cup A||} \\ &= \tau - \frac{||\bar{A} \setminus \bar{A}_{cover}|| - ||\bar{A}_{cover} \setminus \bar{A}||}{||\bar{A} \cup A||} \geq 0. \end{aligned} \quad (3)$$

Second line can be derived from the first line because $(\bar{A}_{cover} \cup A) \subseteq (\bar{A} \cup A)$. In these notions \bar{A}_{cover} and \bar{A} are the the resulting superposition of partial sensing and equivalently the corresponding assemblies of all cells represented in \bar{I}_{cover} and \bar{I}^i , respectively. The search of the subtrees rooted at $I_{i,j} \notin \bar{I}_{cover}$ are terminated, and the next level set $\bar{I}^{i+1} := \bar{I}_{cover}$.

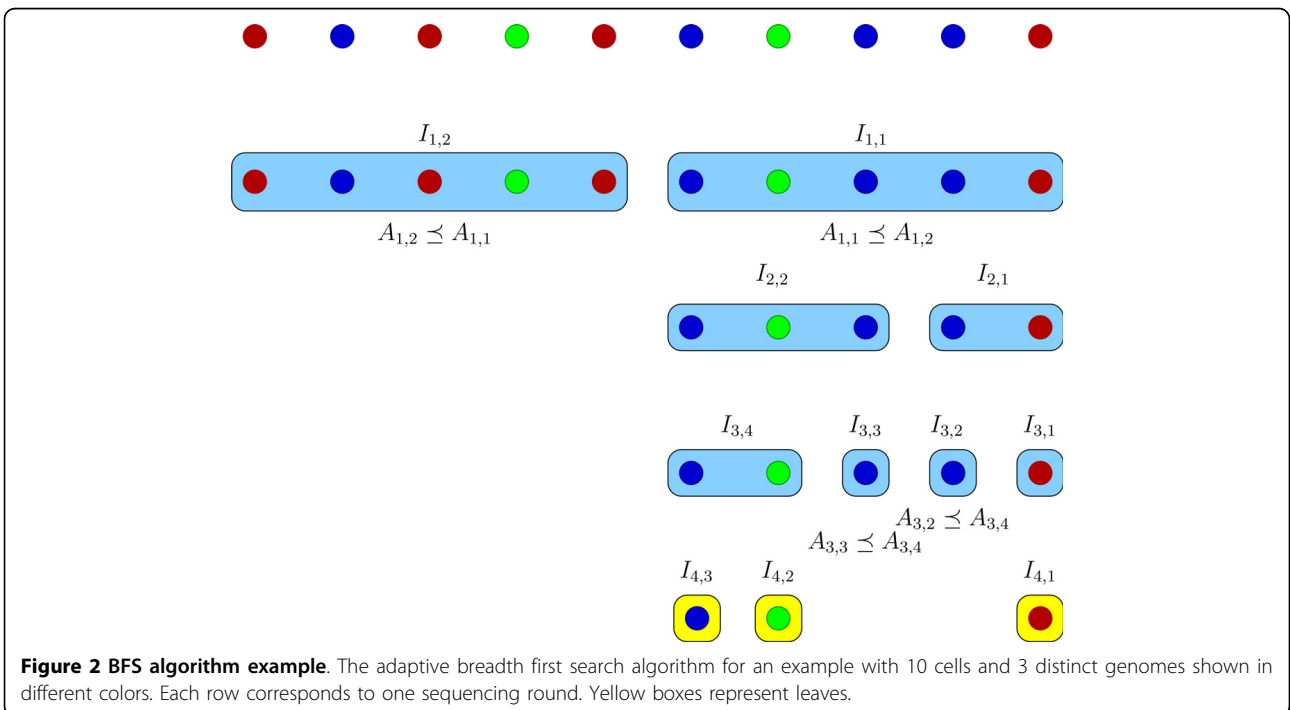
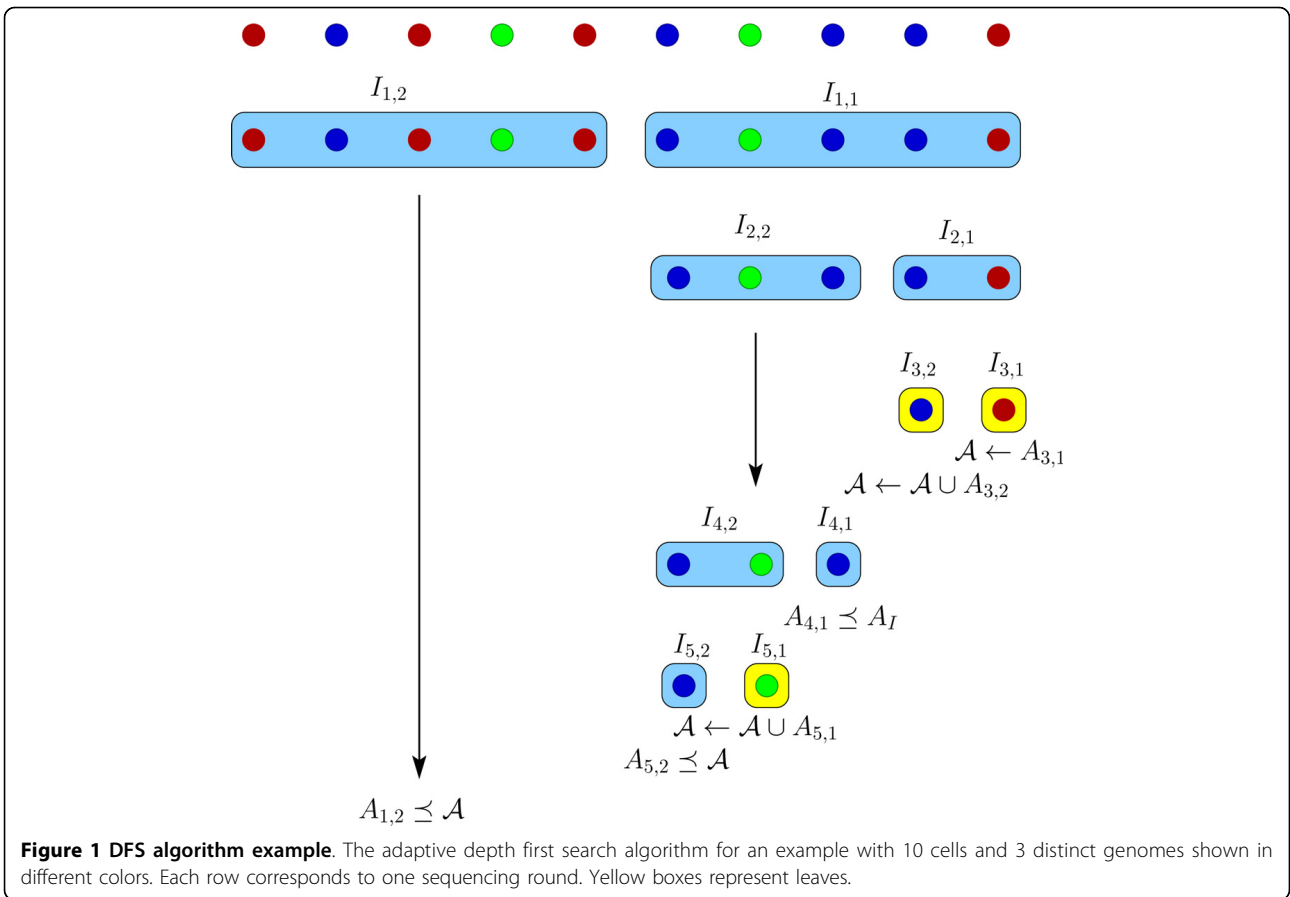
For the DFS strategy, the minimum set cover is calculated gradually during several iterations. Since in the DFS, in each iteration \bar{I}^i only includes two subsets, and the number of cells in both subsets are (almost) equal, the minimum set cover can be calculated based on the greedy algorithm. The $I_{i,j}$ with maximum assembly size has the highest priority to be in the minimum set cover. Therefore, $\bar{I}_{cover} := \{I_{i,j}\}$, and the second $I_{i,j}'$ will be pushed to the stack \bar{W} , which is the waiting list of the untraversed nodes in the tree. If A_{cover} is subsumed in A , then A_{cover} will be emptied and the last element will be popped from \bar{W} and pushed to \bar{I}_{cover} . This will continue until A_{cover} is not subsumed in A . In the end, the next level set $\bar{I}^{i+1} := \bar{I}_{cover}$.

For both search strategies, all subsets in \bar{I}^{i+1} will be divided to two almost equal size subsets, which concludes iteration i . This algorithm will continue until \bar{I}^i and \bar{W} are empty. Figures 1 and 2 depict examples of the DFS and BFS strategies on 10 cells with 3 distinct genomes shown in different colors.

Resource allocation

Resource allocation policy determines the size of partial sensing from each cell in each step. This is done with two objectives: (i) the amount of sensing from each element is such that with a given probability all of the distinct genomes present in $I_{i,j}$ can be reconstructed almost completely from the superposition of partial sensing, and (ii) the total sensing size in the whole algorithm is minimized.

Assume the input set of cells $I_{i,j}$ is obtained from splitting $I_{i-1,k}$ (for clarity we call it I_{parent}). Let t_{parent} , a_{parent} and c_{parent} be total nucleotides sampled, the assembly size, and the average coverage of I_{parent} and $c'_{i,j}$ and $a'_{i,j}$ be the intended coverage and assembly size of $I_{i,j}$. We



assume that there is a constant minimum coverage M_u , such that if the coverage is above M_u , then the resulting assembly covers the entire genome, i.e., does not have any gaps. We would like the actual coverage $c_{i,j}$ after the sequencing and assembly to be at least M_u , so we let $c'_{i,j} = M_u$ as a surrogate. Hence, the total nucleotides $t_{i,j}$ to be sampled and sequenced from $I_{i,j}$ is estimated by

$$t_{i,j} = \frac{t_{parent} c'_{i,j} a'_{i,j}}{c_{parent} a_{parent}} = M_u \frac{t_{parent} a'_{i,j}}{c_{parent} a_{parent}}. \quad (4)$$

In this equation a' should be estimated from c_{parent} and a_{parent} and may differ from the actual value a obtained after sequencing and assembly. If $I_{i,j}$ is a leaf, i.e., $|I_{i,j}| = 1$, then the algorithm does a deep sequencing of the single cell in $I_{i,j}$. In that case, the algorithm repeats the resource allocation and sequencing until a sufficient actual coverage is reached, i.e., $c_{i,j} \geq M_u$.

DFS versus BFS

Although the two search strategies are similar, they have differences in several aspects:

- In DFS, the number of cell subcollections to be explored in each round is fixed. For instance, it is two in the current implementation. In BFS, this number is dynamic. For instance, in the first round of the example in Figure 2, it is two and in the third round is four. From another point of view, the number of rounds in BFS is fixed, $\lceil \log_2 n \rceil$, whereas that in DFS is variable depending on the setup. The number of rounds and the number of subcollections are both desired to be minimum because each incurs a cost: each round incurs a setup cost for sample preparation and running-and-stopping the sequencer and each subcollection requires a unique barcode and incurs the cost of barcoding and sequencing the base pairs in the barcode. Hence, there is a trade-off between the two costs that determines the most suitable algorithm. The optimal algorithm should consider both costs.
- In the process of choosing the minimum set cover, sets are compared to determine the subsumption relationships. In DFS, one side of the comparison is always A in which each single cell is deeply sequenced and completely assembled. Therefore, our information about that side is almost complete and has minimal error. The resource allocation of the algorithm is applied such that $A_{i,j}$'s for all i and j are completely sequenced on average to reduce unwanted missing of some distinct genomes. If the resulting coverage of $A_{i,j}$ is low, i.e., $c_{i,j} \leq M_i$ for a constant M_i , then the data is considered unreliable and $I_{i,j}$ is treated as if it contains a new distinct genome. In that case, $I_{i,j}$ would be

divided into two groups and explored further in the following rounds with increased requested resources. Although $A_{i,j}$'s have deep coverage on average, there is a non-zero probability of missing small and less abundant distinct genomes. We explore these probabilities in the Results section. In BFS, the missing probability is more since on both sides of the \leq relation, cells are deeply sequenced not individually but on average, which increases the error and intrinsic noise in comparisons.

The mentioned differences between the two strategies, do not result in considerable performance priority of one over the other method. Each are proper for a specific condition. In this paper, we compare the two methods using case-by-case setups. We show that the total nucleotides required is almost the same in both methods. We could not do the ensemble analysis for BFS using dynamic programming since the exploration of different subcollections in one round are coupled. For instance, the four subcollections to be compared in the third round of Figure 2 come from two parent subcollections in the second round. That creates an inevitable coupling between the parents in the second round. Without dynamic programming, exploring all of the permutations of cells to provide ensemble analysis is intractable. Therefore, we provide the ensemble analysis only for the DFS algorithm.

Implementation

The pseudocode of the algorithm is given in 1, 2, and 3. COMPRESSEDSEARCH is the main function and SELECTNEXTLEVELSETS and Subsumed are two subfunctions of the algorithm. This algorithm has been implemented in the tool *squeezambler* 2.0.

squeezambler 2.0 versus *squeezambler* 1.0

The tool *squeezambler* 1.0 has been implemented based on the BFS algorithm given in [19]. There are three main differences between the BFS algorithm implemented in *squeezambler* 1.0 and the one in *squeezambler* 2.0:

- In the recursion, the method to choose subsets passed to the next level is different in the two implementations. In *squeezambler* 1.0, every subset that is subsumed in another one is eliminated from further analysis. However, this is not the optimum method to choose next level sets. In *squeezambler* 2.0, a collection of subsets is chosen, which will cover the whole assembly with minimum number of cells.
- The resource allocation in *squeezambler* 1.0 was design for those sequencing technologies that have non-uniform coverage. That resource allocation results in assembly gaps and in some cases causes missing some of the distinct genomes. The resource allocation

in **squeezambler** 2.0 is modified such that it reduces the random missing of the distinct genomes and let us predict the probability of missing genomes. This probability is analysed in the Results section. We have added the resource allocation described in this paper to **squeezambler** 1.0. The new version is called **squeezambler** 1.1.

- The parameter τ in **squeezambler** 1.0 is variable and is dependent on the number of cells involved in each round. As the number of cells increases, τ decreases. Reduced τ increases the number of base pairs required to be sequenced when error appears in the reads. In **squeezambler** 2.0 τ is set to be fixed in the whole algorithm.

Ensemble analysis

We provide the ensemble analysis on the DFS algorithm to calculate the expected total number of sequenced base pairs and the probability of capturing every distinct genome over the entire ensemble of $n!$ permutations of C . To reduce the complexity, instead of exhaustively trying multiple permutations we developed a dynamic programming algorithm to calculate the results. In this analysis, the simulation of the entire process is replaced by black boxes which are mathematical models of the behaviors of the process. To ease finding decoupled effects of different parameters in the algorithm, we do not consider the sequencing and assembly errors in our model. Another assumption, again in the interest of other important sparsity-related parameters, is to consider uniformity of coverage. This assumption is not far from reality. With the advancement of automated microfluidic cell separation and cultivation devices [2,3], the genome can be captured from cultivated cells and sequenced with close to uniform coverage. This is different from the assumption we made in [15] for which a genome was amplified from a single cell using multiple displacement amplification and suffered from highly uneven coverage after sequencing. Although we are assuming uniform coverage distribution in this work as opposed to in [15], this is only a convenient choice that does not change the algorithm. This assumption is reflected only in (5).

Given the uniformity of coverage, we assume

$$a(c) = \begin{cases} g \frac{c}{M_u} & c \leq M_u, \\ g & c \leq M_u, \end{cases} \quad (5)$$

in which c is the sequencing coverage, a is the total assembly size, g is the genome size, and M_u is a constant that defines the minimum coverage to obtain a complete assembly of the whole genome. For more advanced models, see [20]. In the ensemble analysis, we treat (co-) assembly as a black box oracle that knows g and M_u , the

input to which is the total sequenced nucleotides and the output of which are a and c . That is based on the assumption that ac is the total sequenced nucleotides, i.e. there are no sequencing errors. In this case the resource allocation formula in (4) will be reduced to $t_{i,j} = M_u a'_{i,j}$. In the worst case, all of the distinct genomes in I_{parent} are also represented in $I_{i,j}$ which means $a'_{i,j} \geq a_{parent}$. Therefore, the total nucleotides can be estimated by

$$t_{i,j} = 2M_u a_{parent} \quad (6)$$

which is twice the lower bound as a safe margin.

Dynamic programming algorithm

The dynamic programming algorithm can be divided into three main functions Cost (Algorithm 4), ALLOCATESEQUENCEASSEMBLEORACLE (Algorithm 5), and Subsumed (Algorithm 6). Algorithm 4 is the main dynamic programming, and its subroutines are presented in Algorithms 5 and 3. Let s be the number of distinct genomes in C . A distinct genome profile $p = (p_1, p_2, \dots, p_s) \in (\mathbb{N} \cup \{0\})^s$ is a population vector. In the root of the search tree, $\sum_{j=1}^s p_j = n$, where n is the total number of cells. The vector of deeply sequenced and assembled distinct genomes before exploration of the current node is denoted by $I = (I_1, I_2, \dots, I_s)$ where $I_j \in \{0, 1\}$ and $I_j = 1$ means that the distinct genome j has been sequenced and completely assembled. Throughout the algorithms, $\bar{a} = (a_1, a_2, \dots, a_s)$ is the assembly size profile per distinct genome in the current node, and $\|\bar{a}\|_1$ is the total assembly size. Denote the assembly size of the parent search node by a_{parent} , total sequenced nucleotides by t , and the expected total sequenced nucleotides by $E[t]$. We denote the probabilities of capturing distinct genomes by $P = (q_1, q_2, \dots, q_{2^s})$, in which q_j is the probability of the vector of deeply sequenced and assembled distinct genomes $I \in \{0, 1\}^s$ where $j = I((2^0, 2^1, \dots, 2^{s-1})) + 1$, upon complete exploration of the current node in the search tree. Above, j is the decimal representation of I treated as binary (in reverse order) plus one. For example, for $s = 3$ and $I = (0, 1, 1)$ (in short $I = 011$), j is the decimal representation of reverse of 011 plus one which is equal to 7. Therefore, q_7 (or in another notation q_{011}) is the probability that after exploration of the current node the distinct genomes 2 and 3 are recognized by the algorithm, while distinct genome 1 is missed.

The COST function requires the genome population profile for the current node, p , the set of distinct genomes already deeply sequenced and assembled, I , and the result of the assembly of the parent node a_{parent} as input parameters. The output of this function is the estimated cost $E[t]$ and the probabilities of capturing genomes P

after exploring the current node. At first (line 5), an oracle will decide on the sampling size of the current node and the resulting coverage and assembly size based on the formulations given in (4) and (5). In line 7, using the function SUBSUMED, the node is then compared with I to see if it includes any new distinct genome. If there is no new distinct genome (line 41), the probability of capturing the corresponding genomes is set to 1 and the function exits. Otherwise, the node will be explored further.

If the node includes only one cell, then that node is a leaf, and it will be sequenced and assembled deeply (line 9) and the corresponding capturing probability is set to 1. In the case of a node with more than one cell, the collection of cells will be divided into two groups. The expected cost and capturing probability, starting from the current node, is calculated over the ensemble of all of the possible divisions of the node into v and w between lines 19 and 39. The ensemble parameters are averaged over all divisions (v , w) by calculating their probability of occurrence (line 36). For a given division, v is explored followed by w . For each $u \in \{v, w\}$ (line 24) and each combination of already captured distinct genomes with non-zero probability (lines 26-28), the expected cost t' and capturing probability profile P'' are recursively calculated using the COST function (line 29). These parameters are averaged over all non-zero probability profiles (lines 30, 31).

Results

DFS versus BFS

To compare the performance of DFS and BFS algorithms, we tested the algorithms on simulated data. We have selected 6 distinct genomes from human gut microbiome [4] to generate 3 setups (see Tables 1 and 2). The genomes were amplified and sequenced using ART Illumina sequencing simulator [21]. Reads are 100 bp long, uniformly covering the whole genome. The assembler used in the paper is HyDA co-assembler [16]. To allocate resources, the relationship between the coverage and the assembly size of this setup is calculated using ART and HyDA over the 6 distinct genomes. The result is depicted in Figure 3. We selected $M_u = M_l = 15$ and $\tau = 0.1$.

We have compared the performance of four different methods: (i) naïve method of sequencing each cell deeply,

Table 2 Our simulation setups: (i) 31 cells; 6 distinct genomes, (ii) 59 cells; 4 distinct genomes, and (iii) 140 cells; 4 distinct genomes.

NCBI ID	Abundance (%)					
	31 cells; 6 distinct genomes		59 cells; 4 distinct genomes		140 cells; 4 distinct genomes	
NC 004663.1	11	35.50%	22	37%	35	25%
NC 009614.1	4	13%	7	12%	35	25%
NC 009615.1	3	10%	8	14%	35	25%
NC 008532.1	1	3%	0	0%	0	0%
NC 016776.1	1	3%	0	0%	0	0%
FP929051.1	11	35.50%	22	37%	35	25%

(ii) BFS compressed method using **squeezambler** 1.1, (iii) BFS compressed method using **squeezambler** 2.0, and (iv) DFS compressed method using **squeezambler** 2.0. The results are summarized in Table 3. See that **squeezambler** 2.0 BFS outperforms **squeezambler** 1.1 BFS in total sequencing base pairs. The BFS and DFS algorithms of **squeezambler** 2.0 have close performances in terms of the total sequenced base pairs. However, for the case of 31 cells and 6 distinct genomes, BFS missed one distinct genome. This is one of the examples that shows the genome missing probability is slightly more in BFS. Overall, both algorithms have comparable performances. Comparison of the performance of naïve algorithm and compressed methods shows that as the number of cells increases, the total sequenced base pairs increases linearly for the naïve algorithm and sublinear for the compressed methods. Although for the case of 31 cells, the naïve method outperforms the compressed method, at 140 cells the compressed method shows its strength.

Ensemble analysis

In this section, we selected $M_u = 5$, $M_l = 0.3$, and $\tau = 0.2$ except in Table 5. Bacterial genome sizes were considered to be within 1 - 12 Mbps. In the current version of the program implemented in MATLAB, we computed the results for a small number of cells and distinct genomes, i.e., < 200 cells and < 11 distinct genomes. However, we expect to be able to run the algorithm for a larger number of cells and distinct genomes with an

Table 1 6 distinct genomes (species) used in simulation.

NCBI ID	Name	Ref. Status	Size (bps)
NC 004663.1	Bacteroides thetaiotaomicron VPI-5482 chromosome	complete	6.29 M
NC 009614.1	Bacteroides vulgatus ATCC 8482 chromosome	complete	5.16 M
NC 009615.1	Parabacteroides distasonis ATCC 8503 chromosome	complete	4.81 M
NC 008532.1	Streptococcus thermophilus LMD-9	complete	1.86 M
NC 016776.1	Bacteroides fragilis 638R	complete	5.37 M
FP929051.1	Ruminococcus bromii L2-63	draft	2.25 M




Figure 3 Output assembly size in percentile versus coverage.

The plot is the average of the output of sequencing simulation using ART [21] and assembly using HyDA [16] over 6 distinct genomes listed in Table 1.

advanced implementation in C++. We tried to decouple the effect of different parameters in the analysis, namely τ in the algorithm, and the number of cells and species in the input. We would like to test whether the expected total cost complexity is $O(s)$ for a fixed n and $O(\log n)$ for a fixed s and population profile. This is a first step to show the expected total cost is $O(s \log n)$ in the future. Therefore, the results also provide intuition for a potential thorough theoretical analysis of the expected cost and capturing probability.

Expected cost

To investigate the growth rate of $E[t]$ for different number of cells n and compare it with the naïve cell-by-cell sequencing, we ran our program for the profiles of $\frac{n}{8}(3, 2, 3)$ where $n = 8, \dots, 192$ is the total number of

cells with genome sizes (4, 12, 2) Mbps. Figure 4 depicts the results. The total sequenced nucleotides in the naïve case is $M_u \times \text{max genome size} \times n = 60n$ Mbps. The genome of length 2 Mbps may not be captured with at most 3% probability; the other two genomes are always captured. Figure 4 suggests that $E[t]$ grows almost linearly with $\log n$ whereas the naïve cost grows linearly with n . Hence, $E[t] = O(\log n)$ for fixed number of distinct genomes s and population profile.

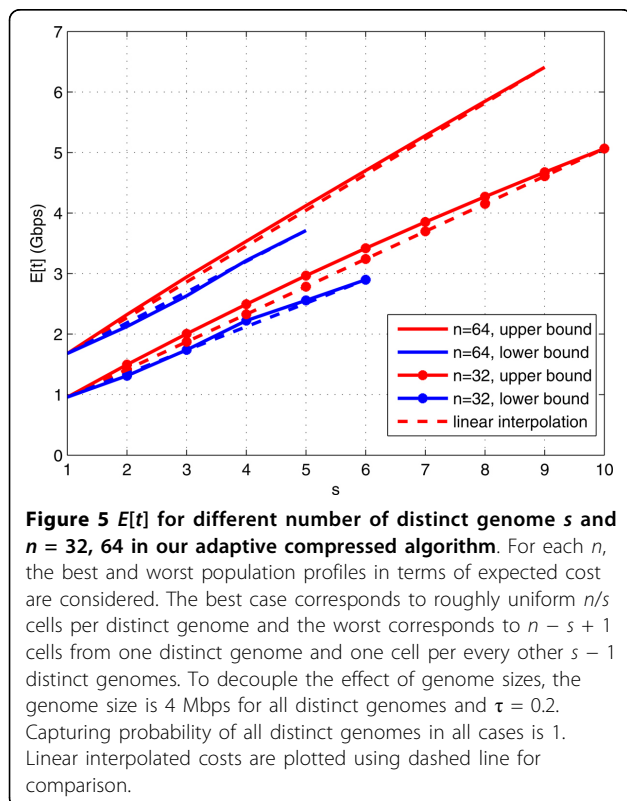
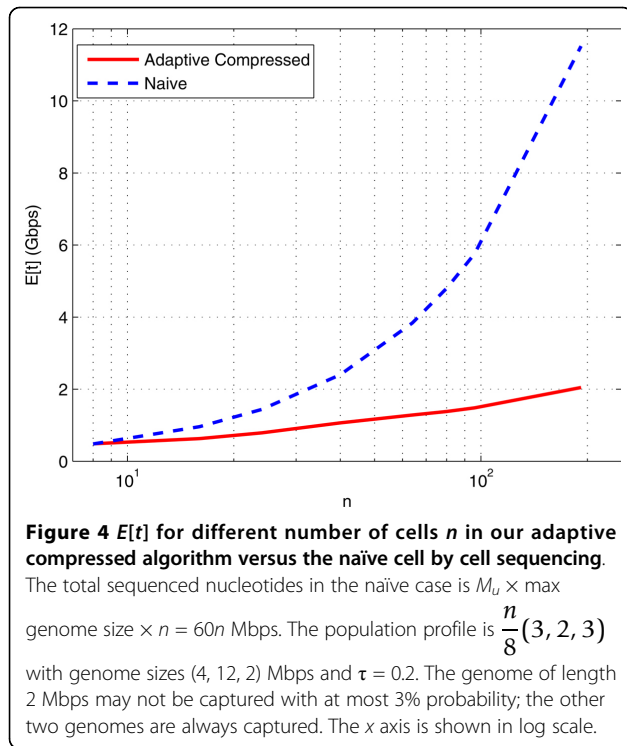
To characterize the behavior of $E[t]$ for different number of distinct genomes s , we plotted $E[t]$ versus $s = 1, \dots, 10$ for $n = 32, 64$ in Figure 5. For each n , the best and worst population profiles in terms of expected cost were considered. The best case corresponds to roughly uniform n/s cells per distinct genome and the worst corresponds to $n - s + 1$ cells from one distinct genome and one cell per every other $s - 1$ distinct genomes (experimentally verified). The genome size was fixed at 4 Mbps for all distinct genomes to factor out the effect of genome sizes and $\tau = 0.2$. Capturing probability of all distinct genomes in all cases was 1. Comparing the experimental curves with the linearly interpolated cost curves in Figure 5 suggests that the upper bound of cost (worst case) for each fixed n is $O(s)$. In other words, the worst case cost grows almost linearly with the number of distinct genomes. Hence, for all population profiles, $E[t] = O(s)$ for fixed number of cells n .

Capturing probability

We ran our program for a number of setups with 3 and 4 distinct genomes. The number of cells n was fixed at 40 in all experiments. Genome sizes varied between 1 and 12 Mbps. The total naïve single-cell sequencing is the minimum coverage M_u times the number of cells, which is 40, times the maximum genome size. In this case, the naïve total cost is 2.4 Gbps. Tables 4 and 5 show the setups and their expected

Table 3 squeezambler results for the three setups summarized in Table 2.

Setup	Method	Total sequencing (Gbps)	Max barcodes	No. of predicted distinct genomes	Iterations
31 cells; 6 distinct genomes	Naïve	4	31	6	1
	squeezambler 1.1, BFS	13.8	10	6	5
	squeezambler 2.0, BFS	11.2	8	5	5
	squeezambler 2.0, DFS	11	2	6	12
59 cells; 4 distinct genomes	naïve	8	59	4	1
	squeezambler 1.1, BFS	10	4	4	6
	squeezambler 2.0, BFS	8.7	4	4	6
	squeezambler 2.0, DFS	8.8	2	4	9
140 cells; 4 distinct genomes	naïve	18.5	140	4	1
	squeezambler 2.0, DFS	10.3	2	4	10



sequenced nucleotides $E[t]$ and P , the capturing probability. Recall $P = (q_{00\dots 0}, q_{10\dots 0}, \dots, q_{11\dots 1})$, in which q_I for $I \in \{0, 1\}^s$ is the joint probability of capturing distinct genome j if $I_j = 1$ and missing it if $I_j = 0$ for $j = 1, 2, \dots, s$.

In Table 4 with constant $\tau = 0.2$, that genome whose length is 2 Mbps in the first row may not be captured with probability 0.62% because $2/12 < 0.2 = \tau$. We see a similar effect in the subsequent rows. Those genomes whose lengths are 1 and 2 Mbps may not be captured with non-zero probability. As the abundances of short genomes increase, the probabilities of missing them decrease. This suggests that those genomes whose lengths are shorter than τ times the largest genome size may not be captured with non-zero probability.

To further investigate, we ran our program on other setups with varying τ , which are presented in Table 5. For the profile (1, 1, 2, 4) and $\tau = 0, 0.1$, all of the genomes are captured with probability 1. For $\tau = 0.2$, there is a non-zero probability of missing the genomes of length 1 and 2 Mbps. For $\tau = 0.4$, the genome of size 4 Mbps joins the other two short genomes, and the probability of missing it becomes non-zero. Comparing the profiles (1, 1, 2, 4) and (1, 1, 1, 5) for $\tau = 0.4$, the probability of missing the genome of size 2 Mbps, i.e., q_{1101} , increases significantly from 0.8% to 23% as a result of decrease in its abundance, whereas the probability of not capturing the genome of size 1 Mbps, i.e., q_{1110} , decreases from 4.1% to 0.8% as its abundance increases. This suggests that the missing probability depends on abundance, potentially relative to the abundance of the largest genome in the population.

Conclusion

We presented an adaptive compressed algorithm for sequencing and *de novo* assembly of distinct genomes in a bacterial community. We used the characteristics of sparsity of distinct genomes in a community of cells to decrease the amount of nucleotides needed to be sequenced. Using a dynamic programming algorithm to analyze the ensemble behavior of the algorithm, we showed that the expected cost is $O(\log(\text{number of cells in the community}))$ for fixed genome population profiles and $O(\text{number of distinct genomes})$ for fixed number of cells. Furthermore, we showed that for a non-zero threshold τ , those genomes whose sizes relative to the maximum genome size in the community is less than τ may go undetected with a non-zero probability. This probability depends on the abundance of the corresponding genome. These results shed light on our future path towards theoretical analysis of our algorithm and further tree exploration strategies.

Table 4 Effect of population profile.

$p/5$	Genome size (Mbps)	$E[t]$ (Gbps)	$P = (q000, q100, q010, q110, q001, q101, q011, q111)$
(3, 2, 3)	(4, 12, 2)	1.066	(0, 0, 0, 0.0062, 0, 0, 0, 0.9938)
			$P = (q0000, q1000, q0100, q1100, q0010, q1010, q0110, q1110, q0001, q1001, q0101, q1101, q0011, q1011, q0111, q1111)$
(1, 1, 1, 5)	(4, 12, 2, 1)	1.383	(0, 0, 0, 0, 0, 0, 0.0020, 0, 0, 0, 0.0271, 0, 0, 0, 0.9709)
(1, 1, 2, 4)	(4, 12, 2, 1)	1.407	(0, 0, 0, 0, 0, 0, 0, 0.0123, 0, 0, 0, 0.0001, 0, 0, 0, 0.9876)
(1, 2, 3, 2)	(4, 12, 2, 1)	1.535	(0, 0, 0, ϵ , 0, 0, 0, 0.1687, 0, 0, 0, 0.0135, 0, 0, 0, 0.8179)

In this setup $\tau = 0.2$ and the total number of cells is $n = 40$. q_l for $l \in \{0, 1\}^s$ is the joint probability of capturing distinct genome j if $l_j = 1$ and missing it if $l_j = 0$ for $j = 1, 2, \dots, s$. Different values of l are explicitly given as binary strings in the fourth column. ϵ is a non-zero probability less than 10^{-4} .

Table 5 Effect of threshold τ .

$p/5$	τ	$E[t]$ (Gbps)	$P = (q0000, q1000, q0100, q1100, q0010, q1010, q0110, q1110, q0001, q1001, q0101, q1101, q0011, q1011, q0111, q1111)$
(1, 1, 2, 4)	0	1.428	(0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1)
(1, 1, 2, 4)	0.1	1.423	(0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1)
(1, 1, 2, 4)	0.2	1.407	(0, 0, 0, 0, 0, 0, 0, 0.0123, 0, 0, 0, 0.0001, 0, 0, 0, 0.9876)
(1, 1, 2, 4)	0.4	1.266	(0, 0, 0, 0.0002, 0, 0, 0.0008, 0.0411, 0, 0, ϵ , 0.0794, 0, 0, 0.1621, 0.7165)
(1, 1, 1, 5)	0	1.418	(0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1)
(1, 1, 1, 5)	0.1	1.414	(0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1)
(1, 1, 1, 5)	0.2	1.383	(0, 0, 0, 0, 0, 0, 0, 0.0020, 0, 0, 0, 0.0271, 0, 0, 0, 0.9709)
(1, 1, 1, 5)	0.4	1.214	(0, 0, 0, 0.0013, 0, 0, ϵ , 0.0078, 0, 0, ϵ , 0.2308, 0, 0, 0.1369, 0.6231)

The genome sizes are (4, 12, 2, 1) Mbps and the total number of cells is $n = 40$. q_l for $l \in \{0, 1\}^s$ is the joint probability of capturing distinct genome j if $l_j = 1$ and missing it if $l_j = 0$ for $j = 1, 2, \dots, s$. Different values of l are explicitly given as binary strings in the fourth column. ϵ is a non-zero probability less than 10^{-4} .

Algorithms

Algorithm 1 Compressed Sequencing

- 1: **Input:** $C = \{C_1, C_2, \dots, C_n\}$
- 2: **Output:** A, I
- 3: $I_{1,1} \leftarrow \{C_1, \dots, C[n/2]\}$
- 4: $I_{1,2} \leftarrow \{C[n/2]+1, \dots, C_n\}$
- 5: $\bar{I} \leftarrow \{I_{1,1}, I_{1,2}\} \triangleright \bar{I}$ is the list of the subsets to be analysed in the subsequent round
- 6: $\bar{W} \leftarrow \{\}$ $\triangleright \bar{W}$ is the waiting list of the subsets assembled but not ready to be analysed immediately
- 7: $i \leftarrow 1 \triangleright i$ is the sequencing round index
- 8: **while** EITHER \bar{I} OR \bar{W} IS NOT EMPTY **do**
- 9: $\bar{t} \leftarrow \text{RESOURCEALLOCATE}(\bar{I}, a_{parent}, c_{parent}, M_{i,i}) \triangleright \bar{A}, \bar{t} = \{t_{i,1}, \dots, t_{i,|\bar{I}|}\}$; based on Equ 4
- 10: $\bar{A}, \bar{a}, \bar{c} \leftarrow \text{SEQUENCEANDASSEMBLE}(\bar{t}, C) \triangleright \bar{c} = \{c_{i,1}, \dots, c_{i,|\bar{I}|}\}$; $\bar{c} = \{c_{i,1}, \dots, c_{i,|\bar{I}|}\}$; $A_{i,j}, c_{i,j}$ are the assembly set and the average coverage of $I_{i,j}$ respectively.
- 11: $\bar{I}, \bar{W} \leftarrow \text{SELECTNEXTLEVELSETS}(\bar{I}, \bar{W}, \bar{A}, \bar{c}, F) \triangleright F$: DFS or BFS flag
- 12: $i \leftarrow i + 1$
- 13: **end while**
- Algorithm 2** selectNextLevelSets
- 1: **Input:** $\bar{I}, \bar{W}, \bar{A}, \bar{c}, F$
- 2: **Output:** $\bar{I}_{new}, \bar{W}_{new}$
- 3: $\bar{L} = \{\}$ \triangleright list of subsets with low quality assemblies
- 4: $A_L = \{\}$ \triangleright assemblies of subsets in \bar{L}
- 5: **for** $j = 1 \dots |\bar{I}|$ **do**

- 6: **if** $c_{i,j} < M_i$ **then**
- 7: **MOVE** $I_{i,j}$ FROM \bar{I} TO $\bar{L} \triangleright$ move all low coverage assembled $I_{i,j}$ TO \bar{L}
- 8: **MOVE** $A_{i,j}$ FROM \bar{A} TO A_L
- 9: **else**
- 10: **if** $|I_{i,j}| = 1$ **then**
- 11: **MOVE** $I_{i,j}$ FROM \bar{I} TO $I \triangleleft$ move all single cell assembled $I_{i,j}$ TO I
- 12: **MOVE** $A_{i,j}$ FROM \bar{A} TO A
- 13: **end if**
- 14: **end if**
- 15: **end for**
- 16: **if** F is BFS **then**
- 17: **FIND THE MINIMUM SET COVER** A_{new} CORRESPONDING TO $\bar{I}_{new} \subseteq \bar{I}$ FOR WHICH $D_\tau((A \cup A_L \cup A), (\bar{A}_{new} \cup A_L \cup A)) \geq \text{AND } |\bar{I}_{new}|$ IS MINIMUM \triangleright Equ. 3
- 18: $\bar{I}_{new} \leftarrow \bar{I}_{new} \cup \bar{L}$
- 19: $\bar{W}_{new} = \{\}$
- 20: **else**
- 21: **if** F is DFS **then**
- 22: $\bar{W}_{new} = \bar{W}$
- 23: **POP ALL** $I_{i,j}$'s FROM \bar{I}_{new} AND PUSH TO \bar{W}_{new} EXCEPT ONE SUBSET WITH THE MAXIMUM ASSEMBLY SIZE
- 24: **while** $|\bar{I}_{new}| = 0$ AND $|\bar{W}_{new}| > 0$ **do**
- 25: $\bar{I}_{new} \leftarrow$ POP LAST SUBSET IN \bar{W}_{new}
- 26: $A_{new} \leftarrow$ ASSEMBLY OF \bar{I}_{new}

```

27:         if NOT SUBSUMED ( $A_{new}$   $A$ ) then
28:             PUSH  $I_{new}$  TO  $\bar{I}_{new}$ 
29:         end if
30:     end while
31: end if
32: end if
33: DIVIDE ALL  $I_{ij}$ 's IN  $\bar{I}_{new}$  TO TWO SETS
Algorithm 3 SUBSUMED
1: Input:  $A_{new}$   $A$ 
2: Output:  $r \in \{ \text{true}, \text{false} \}$ 
3:      $\triangleright$  low quality assembly; explore the
node further.
4: if  $c \leq M_l$  then
5:      $r \leftarrow \text{false}$ 
6: else
7:      $D \leftarrow \tau - ||A_{new} \setminus A|| / ||A_{new}||$       $\triangleright$  Equ 2
8:     if  $D < 0$  then      $\triangleright$  Equ 1
9:          $r \leftarrow \text{false}$ 
10:    else
11:         $r \leftarrow \text{true}$ 
12:    end if
13: end if
Algorithm 4 COST - ensemble analysis main function
1: Input:  $p, I, a_{parent}$ 
2: Output:  $E[t]$  and  $\mathbb{P}$ 
3:
4:  $\mathbb{P} = (q_1, q_2, \dots, q_{2s}) \leftarrow 0$ 
5:  $t, \bar{a}, c \leftarrow \text{ALLOCATESEQUENCEASSEMBLEORACLE}(p, a_{parent})$   $\triangleright t$  total nt,  $\bar{a} = (a_1, \dots, a_s), c$  coverage
6:  $E[t] \leftarrow t$ 
7: if not SUBSUMEDENSEMBLE( $\bar{a}, c, I$ ) then
8:     if  $||p||_1 = 1$  then  $\triangleright$  leaf base case
9:         while  $c < 2M_u$  do  $\triangleright$  ensures the complete
assembly for leaves
10:             $t, \bar{a}, c \leftarrow \text{ALLOCATESEQUENCEASSEMBLEORACLE}(p, ||\bar{a}||_1)$ 
11:             $E[t] \leftarrow E[t] + t$ 
12:        end while
13:         $k \leftarrow \arg \max p_j$ 
14:         $I^{new} \leftarrow I$ 
15:         $I_k^{new} \leftarrow 1$ 
16:         $j \leftarrow \langle I^{new}, (2^0, 2^1, \dots, 2^{s-1}) \rangle + 1$ 
17:         $q_j = 1$   $\triangleright$  updating  $\mathbb{P}$ 
18:    else  $\triangleright$  recursion
19:        for  $v + w = p, v, w \in (N \cup \{0\})^s$  do
20:             $t \leftarrow 0$ 
21:             $\mathbb{P}' \leftarrow \mathbb{P} \triangleright \mathbb{P}' = (q'_1, q'_2, \dots, q'_{2s})$ 
22:             $j \leftarrow \langle I, (2^0, 2^1, \dots, 2^{s-1}) \rangle + 1$ 
23:             $q'_j \leftarrow 1$ 
24:            for  $u \in \{v, w\}$  do  $\triangleright v$  is explored followed
by  $w$ 
25:                 $\mathbb{P}'' \leftarrow 0$ 
26:                for  $b$  A BINARY VECTOR OF LENGTH  $s$  do

```

```

27:                     $j \leftarrow \langle b, (2^0, 2^1, \dots, 2^{s-1}) \rangle + 1$ 
28:                    if  $q'_j > 0$  then  $\triangleright$  average over all
already captured distinct genome profile with non-zero
probability
29:                         $t', \mathbb{P}''' \leftarrow \text{COST}(u, b, ||\bar{a}||_1)$ 
30:                         $t \leftarrow t + t'q'_j$ 
31:                         $\mathbb{P}'' \leftarrow \mathbb{P}'' + q'_j\mathbb{P}'''$ 
32:                    end if
33:                end for
34:             $\mathbb{P}' \leftarrow \mathbb{P}' \triangleright$  for  $w, q''$ 's are updated
end for
35:        end for
36:         $\pi \leftarrow \prod_{j=1}^s \binom{p_j}{v_j} / \binom{||p||_1}{||v||_1}$   $\triangleright$  probability of
( $v, w$ )
37:         $\mathbb{E}[t] \leftarrow \mathbb{E}[t] + \pi t$   $\triangleright$  average over all possible
( $v, w$ )
38:         $\mathbb{P} \leftarrow \mathbb{P} + \pi\mathbb{P}'$   $\triangleright$  average over all possible
( $v, w$ )
39:    end for
40: end if
41: else  $\triangleright$  all distinct genomes represented in  $p$  have
already been sequenced
42:      $j \leftarrow \langle I, (2^0, 2^1, \dots, 2^{s-1}) \rangle + 1$ 
43:      $q_j = 1$   $\triangleright$  updating  $\mathbb{P}$ 
44: end if
Algorithm 5 ALLOCATESEQUENCEASSEMBLE-
ORACLE - ensemble analysis
1: Input:  $p, a_{parent}$ 
2: Output:  $t, \bar{a}$ , and  $c$ 
3:
4:  $t \leftarrow 2M_{uaparent}$       $\triangleright$  Equ 6
5:  $t_{pc} \leftarrow t / ||p||_1$       $\triangleright$  total sequenced nt per cell
6:  $\bar{t}_{pdg} \leftarrow t_{pc} \cdot p$       $\triangleright$  total sequenced nt per dis-
tinct genome
7:  $\bar{a}, c \leftarrow \text{ORACLE}(\bar{t}_{pdg})$       $\triangleright$  oracle decides on
the assembly size and coverage based on Equ 5
Algorithm 6 SUBSUMEDENSEMBLE - ensemble
analysis
1: Input:  $\bar{a}, c, I$ 
2: Output:  $r \in \{ \text{true}, \text{false} \}$ 
3:
4: if  $c \leq M_l$  then  $\triangleright$  low quality assembly; explore the
node further.
5:      $r \leftarrow \text{false}$ 
6: else
7:      $x \leftarrow (\neg I, \bar{a})$   $\triangleright$  exclusive part of assemblies,  $\neg$ 
is bitwise not, based on Equ 2
8:     if  $\tau - x/\tau - x/||\bar{a}|| < 0$  then  $\triangleright$  Equ 1
9:          $r \leftarrow \text{false}$ 
10:    else
11:         $r \leftarrow \text{true}$ 
12:    end if
13: end if

```

Competing interests

The author declares that they have no competing interests.

Authors' contributions

Z.T. designed and implemented the algorithms, ran the experiments, and wrote the manuscript.

This article has been published as part of BMC Bioinformatics Volume 15 Supplement 9, 2014: Proceedings of the Fourth Annual RECOMB Satellite Workshop on Massively Parallel Sequencing (RECOMB-Seq 2014). The full contents of the supplement are available online at <http://www.biomedcentral.com/bmcbioinformatics/supplements/15/S9>.

Declarations

This work and publication has been supported by Wayne State University.

Published: 10 September 2014

References

- Chitsaz H, Yee-Greenbaum JL, Tesler G, Lombardo M-J, Dupont CL, Badger JH, Novotny M, Rusch DB, Fraser LJ, Gormley NA, Schulz-Trieglaff O, Smith GP, Evers DJ, Pevzner PA, Lasken RS: **Efficient de novo assembly of single-cell bacterial genomes from short-read data sets.** *Nature Biotech* 2011, **29**(10):915-921.
- Zengler K, Toledo G, Rapp'e M, Elkins J, Mathur EJ, Short JM, Keller M: **Cultivating the uncultured.** *Proc Nat Acad Sci* 2002, **99**(24):15681-15686.
- Fitzsimons MS, Novotny M, Lo C-C, Dichosa AE, Yee-Greenbaum JL, Snook JP, Gu W, Chertkov O, Davenport KW, McMurry K, et al: **Nearly finished genomes produced using gel microdroplet culturing reveal substantial intraspecies genomic diversity within the human microbiome.** *Genome research* 2013, **23**(5):878-888.
- Qin J, Li R, Raes J, Arumugam M, Burgdorf KS, Manichanh C, Nielsen T, Pons N, Levenez F, Yamada T, Mende DR, Li J, Xu J, Li S, Li D, Cao J, Wang B, Liang H, Zheng H, Xie Y, Tap J, Lepage P, Bertalan M, Batto JM, Hansen T, Le Paslier D, Linneberg A, Nielsen HB, Pelletier E, Renault P, Sicheritz-Ponten T, Turner K, Zhu H, Yu C, Li S, Jian M, Zhou Y, Li Y, Zhang X, Li S, Qin N, Yang H, Wang J, Brunak S, Dore J, Guarner F, Kristiansen K, Pedersen O, Parkhill J, Weissenbach J, Bork P, Ehrlich SD, Wang J, Antonin M, Artiguenave F, Blottiere H, Borrueil N, Bruls T, Casellas F, Chervaux C, Cultrone A, Delorme C, Denariac G, Dervyn R, Forte M, Friss C, van de Guchte M, Guedon E, Haimet F, Jamet A, Juste C, Kaci G, Kleerebezem M, Knol J, Kristensen M, Layec S, Le Roux K, Leclerc M, Maguin E, Minardi RM, Oozeer R, Rescigno M, Sanchez N, Tims S, Torrejon T, Varela E, de Vos W, Winogradsky Y, Zoetendal E: **A human gut microbial gene catalogue established by metagenomic sequencing.** *Nature* 2010, **464**(7285):59-65.
- Candès EJ, Tao T: **Decoding by linear programming.** *IEEE Transactions on Information Theory* 2005, **51**(12):4203-4215.
- Candès EJ, Tao T: **Near-optimal signal recovery from random projections: Universal encoding strategies?** *IEEE Transactions on Information Theory* 2006, **52**(12):5406-5425.
- Donoho DL: **Compressed sensing.** *Information Theory. IEEE Transactions* 2006, **52**(4):1289-1306.
- Haupt J, Castro RM, Nowak R: **Distilled sensing: Adaptive sampling for sparse detection and estimation.** *IEEE Transactions on Information Theory* 2011, **57**(9):6222-6235.
- Wei D, Hero AO: **Multistage adaptive estimation of sparse signals.** *IEEE Statistical Signal Processing Workshop (SSP)* 2012, 153-156.
- Erlich Y, Gordon A, Brand M, Hannon GJ, Mitra PP: **Compressed genotyping.** *Information Theory. IEEE Transactions* 2010, **56**(2):706-723.
- Stobbe P, Krause A: **Learning fourier sparse set functions.** *Journal of Machine Learning Research - Proceedings Track* 2012, **22**:1125-1133.
- Malloy ML, Nowak RD: **Near-optimal adaptive compressed sensing.** *Signals, Systems and Computers (ASILOMAR)* 2012, 1935-1939, 2012 Conference Record of the Forty Sixth Asilomar Conference On IEEE.
- Kane DM, Nelson J, Woodruff DP: **Proceedings of the Twenty-ninth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems. An optimal algorithm for the distinct elements problem.** 2010, 41-52, ACM.
- Efron B, Thisted R: **Estimating the number of unseen species: How many words did shakespeare know?** 1976, **63**(3):435-447.
- Taghavi Z, Movahedi NS, Drăghici S, Chitsaz H: **Distilled single-cell genome sequencing and de novo assembly for sparse microbial communities.** *Bioinformatics* 2013, **29**(19):2395-2401.
- Movahedi NS, Forouzmand E, Chitsaz H: **De novo co-assembly of bacterial genomes from multiple single cells.** *IEEE Conference on Bioinformatics and Biomedicine* 2012, 561-565.
- Compeau PEC, Pevzner PA, Tesler G: **How to apply de bruijn graphs to genome assembly.** *Nature biotechnology* 2011, **29**(11):987-991.
- Johnson WB, Lindenstrauss J: **Extensions of lipschitz mappings into a hilbert space.** *Contemporary mathematics* 1984, **26**(189-206):1.
- Taghavi Z, Draghici S: **MDAsim: A multiple displacement amplification simulator.** *IEEE Conference on Bioinformatics and Biomedicine* 2012, 575-578.
- Motahari AS, Bresler G, Tse DNC: **Information theory of DNA shotgun sequencing.** *Information Theory IEEE Transactions on* 2013, **59**(10):6273-6289.
- Huang W, Li L, Myers JR, Marth GT: **ART: a next-generation sequencing read simulator.** *Bioinformatics* 2012, **28**(4):593-594.

doi:10.1186/1471-2105-15-S9-S13

Cite this article as: Taghavi: Ensemble analysis of adaptive compressed genome sequencing strategies. *BMC Bioinformatics* 2014 **15**(Suppl 9):S13.

Submit your next manuscript to BioMed Central and take full advantage of:

- Convenient online submission
- Thorough peer review
- No space constraints or color figure charges
- Immediate publication on acceptance
- Inclusion in PubMed, CAS, Scopus and Google Scholar
- Research which is freely available for redistribution

Submit your manuscript at
www.biomedcentral.com/submit

